To Reach or not to Reach? Efficient Algorithms for Total-Payoff Games*

Thomas Brihaye[1], Gilles Geeraerts[2], Axel Haddad[1], and Benjamin Monmege[2]

1 **Université de Mons, Belgium,** thomas.brihaye,axel.haddad@umons.ac.be
2 **Université libre de Bruxelles, Belgium,** gigeerae,benjamin.monmege@ulb.ac.be

arXiv:1407.5030v4 [cs.GT] 14 Jul 2015

Abstract

Quantitative games are two-player zero-sum games played on directed weighted graphs. Total-payoff games—that can be seen as a refinement of the well-studied mean-payoff games—are the variant where the payoff of a play is computed as the sum of the weights. Our aim is to describe the first pseudo-polynomial time algorithm for total-payoff games in the presence of arbitrary weights. It consists of a non-trivial application of the value iteration paradigm. Indeed, it requires to study, as a milestone, a refinement of these games, called min-cost reachability games, where we add a reachability objective to one of the players. For these games, we give an efficient value iteration algorithm to compute the values and optimal strategies (when they exist), that runs in pseudo-polynomial time. We also propose heuristics to speed up the computations.

1998 ACM Subject Classification D.2.4 Software/Program Verification, F.3.1 Specifying and Verifying and Reasoning about Programs

Keywords and phrases Games on graphs; Reachability; Quantitative games; Value iteration

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

Games played on graphs are nowadays a well-studied and well-established model for the computer-aided design of computer systems, as they enable *automatic synthesis* of systems that are *correct-by-construction*. Of particular interest are *quantitative games*, that allow one to model precisely *quantitative* parameters of the system, such as energy consumption. In this setting, the game is played by two players on a directed weighted graph, where the edge weights model, for instance, a cost or a reward associated to the moves of the players. Each vertex of the graph belongs to one of the two players who compete by moving a token along the graph edges, thereby forming an infinite path called a *play*. With each play is associated a real-valued *payoff* computed from the sequence of edge weights along the play. The traditional payoffs that have been considered in the literature include total-payoff [10], mean-payoff [7] and discounted-payoff [17]. In this quantitative setting, one player aims at maximising the payoff while the other tries to minimise it. So one wants to compute, for each player, the best payoff that he can guarantee from each vertex, and the associated optimal strategies (i.e., that guarantee the optimal payoff no matter how the adversary is playing).

Such quantitative games have been extensively studied in the literature. Their associated decision problems (*is the value of a given vertex above a given threshold?*) are known to be

* The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under Grant Agreement n°601148 (CASSTING).

© Thomas Brihaye, Gilles Geeraerts, Axel Haddad and Benjamin Monmege;
licensed under Creative Commons License CC-BY
Conference title on which this volume is based on.
Editors: Billy Editor and Bill Editors; pp. 1–35

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

in NP ∩ co-NP. Mean-payoff games have arguably been best studied from the algorithmic point of view. A landmark is Zwick and Paterson's pseudo-polynomial time (i.e., polynomial in the weighted graph when weights are encoded in unary) algorithm [17], using the *value iteration* paradigm that consists in computing a sequence of vectors of values that converges towards the optimal values of the vertices. After a fixed, pseudo-polynomial, number of steps, the computed values are precise enough to deduce the actual values of all vertices. Better pseudo-polynomial time algorithms have later been proposed, e.g., in [1, 4, 6], also achieving sub-exponential expected running time by means of randomisation.

In this paper, we focus on *total-payoff games*. Given an infinite play π, we denote by $\pi[k]$ the prefix of π of length k, and by $\mathbf{TP}(\pi[k])$ the (finite) sum of all edge weights along this prefix. The *total-payoff* of π, $\mathbf{TP}(\pi)$, is the inferior limit of all those sums, i.e., $\mathbf{TP}(\pi) = \liminf_{k\to\infty} \mathbf{TP}(\pi[k])$. Compared to mean-payoff (and discounted-payoff) games, the literature on total-payoff games is less extensive. Gimbert and Zielonka have shown [10] that optimal memoryless strategies always exist for both players and the best algorithm to compute the values runs in exponential time [9], and consists in iteratively improving strategies. Other related works include *energy games* where one player tries to optimise its energy consumption (computed again as a sum), keeping the energy level always above 0 (which makes difficult to apply techniques solving those games in the case of total-payoff); and a probabilistic variant of total-payoff games, where the weights are restricted to be non-negative [5]. Yet, we argue that the total-payoff objective is interesting as a *refinement* of the mean-payoff. Indeed, recall first that the total-payoff is finite if and only if the mean-payoff is null. Then, the computation of the total-payoff enables a finer, two-stage analysis of a game $\mathcal{G}$: (*i*) compute the mean payoff $\mathbf{MP}(\mathcal{G})$; (*ii*) subtract $\mathbf{MP}(\mathcal{G})$ from all edge weights, and scale the resulting weights if necessary to obtain integers. At that point, one has obtained a new game $\mathcal{G}'$ with null mean-payoff; (*iii*) compute $\mathbf{TP}(\mathcal{G}')$ to *quantify the amount of fluctuation around the mean-payoff* of the original game. Unfortunately, so far, no efficient (i.e., pseudo-polynomial time) algorithms for total-payoff games have been proposed, and straightforward adaptations of Zwick and Paterson's value iteration algorithm for mean-payoff do not work, as we demonstrate at the end of Section 2. In the present article, we fill in this gap by introducing the first pseudo-polynomial time algorithm for computing the values in total-payoff games.

Our solution is a non-trivial value iteration algorithm that proceeds through nested fixed points (see Algorithm 2). A play of a total-payoff game is infinite by essence. We transform the game so that one of the players (the minimiser) must ensure a *reachability objective*: we assume that the game ends once this reachability objective has been met. The intuition behind this transformation, that stems from the use of an inferior limit in the definition of the total-payoff, is as follows: in any play π whose total-payoff is *finite*, there is a position ℓ in the play after which all the partial sums $\mathbf{TP}(\pi[i])$ (with $i \geqslant \ell$) will be larger than or equal to the total-payoff $\mathbf{TP}(\pi)$ of π, and infinitely often both will be equal. For example, consider the game depicted in Figure 1(a), where the maximiser player (henceforth called Max) plays with the round vertices and the minimiser (Min) with the square vertices. For both players, the optimal value when playing from v_1 is 2, and the play $\pi = v_1v_2v_3\ v_4v_5\ v_4v_3\ (v_4v_5)^\omega$ reaches this value (i.e., $\mathbf{TP}(\pi) = 2$). Moreover, for all $k \geqslant 7$: $\mathbf{TP}(\pi[k]) \geqslant \mathbf{TP}(\pi)$, and infinitely many prefixes ($\pi[8]$, $\pi[10]$, $\pi[12]$, ...) have a total-payoff of 2, as shown in Figure 1(b).

Based on this observation, we transform a total-payoff game $\mathcal{G}$, into a new game that has *the same value as the original total-payoff game* but incorporates a reachability objective for Min. Intuitively, in this new game, we allow a new action for Min: after each play prefix $\pi[k]$, he can ask to *stop the game*, in which case the payoff of the play is the payoff $\mathbf{TP}(\pi[k])$

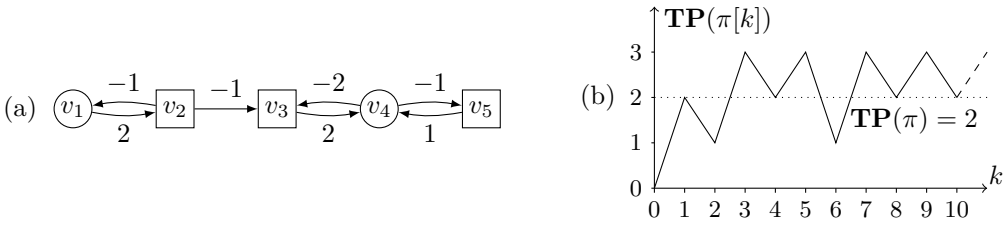

Figure 1 (a) A total-payoff game, and (b) the evolution of the partial sums in π.

of the prefix. However, allowing Min to stop the game at any moment would not allow to obtain the same value as in the original total-payoff game: for instance, in the example of Figure 1(a), Min could secure value 1 by asking to stop after $\pi[2]$, which is strictly smaller that the actual total-payoff (2) of the whole play π. So, we allow Max to *veto* to stop the game, in which case both must go on playing. Again, allowing Max to turn down all of Min's requests would be unfair, so we parametrise the game with a natural number K, which is the maximal number of vetoes that Max can play (and we denote by $\mathcal{G}^K$ the resulting game). For the *play* depicted in Figure 1(b), letting $K = 3$ is sufficient: trying to obtain a better payoff than the optimal, Min could request to stop after $\pi[0]$, $\pi[2]$ and $\pi[6]$, and Max can veto these three requests. After that, Max can safely accept the next request of Min, since the total payoff of all prefixes $\pi[k]$ with $k \geqslant 6$ are larger than or equal to $\mathbf{TP}(\pi) = 2$. Our key technical contribution is to show that *for all total-payoff games, there exists a finite, pseudo-polynomial, value of K such that the values in $\mathcal{G}^K$ and $\mathcal{G}$ coincide* (assuming all values are finite in $\mathcal{G}$: we treat the $+\infty$ and $-\infty$ values separately). Now, assume that, when Max accepts to stop the game (possibly because he has exhausted the maximal number K of vetoes), the game moves to a *target state*, and stops. By doing so, we effectively reduce the computation of the values in the total-payoff game $\mathcal{G}$ to the computation of the values in the total-payoff game $\mathcal{G}^K$ *with an additional reachability objective* (the target state) for Min.

In the following, such refined total-payoff games—where Min *must* reach a designated target vertex—will be called *min-cost reachability games*. Failing to reach the target vertices is the worst situation for Min, so the payoff of all plays that do not reach the target is $+\infty$, irrespective of the weights along the play. Otherwise, the payoff of a play is the sum of the weights up to the first occurrence of the target. As such, this problem nicely generalises the classical shortest path problem in a weighted graph. In the one-player setting (considering the point of view of Min for instance), this problem can be solved in polynomial time by Dijkstra's and Floyd-Warshall's algorithms when the weights are non-negative and arbitrary, respectively. In [11], Khachiyan *et al.* propose an extension of Dijkstra's algorithm to handle the two-player, non-negative weights case. However, in our more general setting (two players, arbitrary weights), this problem has, as far as we know, not been studied as such, except that the associated decision problem is known to be in NP ∩ co-NP [8]. A pseudo-polynomial time algorithm to solve a very close problem, called the *longest shortest path problem* has been introduced by Björklund and Vorobyov [1] to eventually solve mean-payoff games. However, because of this peculiar context of mean-payoff games, their definition of the length of a path differs from our definition of the payoff and their algorithm cannot be easily adapted to solve our min-cost reachability problem. Thus, as a second contribution, we show that a value iteration algorithm enables us to compute in pseudo-polynomial time the values of a min-cost reachability game. We believe that min-cost reachability games bear their own

potential theoretical and practical applications[1]. Those games are discussed in Section 3. In addition to the pseudo-polynomial time algorithm to compute the values, we show how to compute optimal strategies for both players and characterise them: there is always a memoryless strategy for the maximiser player, but we exhibit an example (see Figure 2(*a*)) where the minimiser player needs (finite) memory. Those results on min-cost reachability games are exploited in Section 4 where we introduce and prove correct our efficient algorithm for total-payoff games.

Finally, we briefly present our implementation in Section 5, using as a core the numerical model-checker PRISM. This allows us to describe some heuristics able to improve the practical performances of our algorithms for total-payoff games and min-cost reachability games on certain subclasses of graphs. More technical explanations and full proofs may be found in an extended version of this article [2].

2 Quantitative games with arbitrary weights

We denote by $\mathbb{Z}$ the set of integers, and $\mathbb{Z}_\infty = \mathbb{Z} \cup \{-\infty, +\infty\}$. The set of vectors indexed by V with values in S is denoted by S^V. We let $\preccurlyeq$ be the pointwise order over $\mathbb{Z}_\infty^V$, where $x \preccurlyeq y$ if and only if $x(v) \leqslant y(v)$ for all $v \in V$.

We consider two-player turn-based games on weighted graphs and denote the two *players* by Max and Min. A *weighted graph* is a tuple $\langle V, E, \omega \rangle$ where $V = V_{\mathsf{Max}} \uplus V_{\mathsf{Min}}$ is a finite set of vertices partitioned into the sets V_{Max} and V_{Min} of Max and Min respectively, $E \subseteq V \times V$ is a set of *directed edges*, $\omega \colon E \to \mathbb{Z}$ is the *weight function*, associating an integer weight with each edge. In our drawings, Max vertices are depicted by circles; Min vertices by boxes. For every vertex $v \in V$, the set of successors of v by E is denoted by $E(v) = \{v' \in V \mid (v, v') \in E\}$. Without loss of generality, we assume that every graph is deadlock-free, i.e., for all vertices v, $E(v) \neq \emptyset$. Finally, throughout this article, we let $W = \max_{(v,v') \in E} |\omega(v, v')|$ be the greatest edge weight (in absolute value) in the game graph. A *finite play* is a finite sequence of vertices $\pi = v_0 v_1 \cdots v_k$ such that for all $0 \leqslant i < k$, $(v_i, v_{i+1}) \in E$. A *play* is an infinite sequence of vertices $\pi = v_0 v_1 \cdots$ such that every finite prefix $v_0 \cdots v_k$, denoted by $\pi[k]$, is a finite play.

The total-payoff of a finite play $\pi = v_0 v_1 \cdots v_k$ is obtained by summing up the weights along π, i.e., $\mathbf{TP}(\pi) = \sum_{i=0}^{k-1} \omega(v_i, v_{i+1})$. In the following, we sometimes rely on the mean-payoff to obtain information about total-payoff objectives. The *mean-payoff* computes the average weight of π, i.e., if $k \geqslant 1$, $\mathbf{MP}(\pi) = \frac{1}{k} \sum_{i=0}^{k-1} \omega(v_i, v_{i+1})$, and $\mathbf{MP}(\pi) = 0$ when $k = 0$. These definitions are lifted to infinite plays as follows. The total-payoff of a play π is given by $\mathbf{TP}(\pi) = \liminf_{k \to \infty} \mathbf{TP}(\pi[k])$.[2] Similarly, the mean-payoff of a play π is given by $\mathbf{MP}(\pi) = \liminf_{k \to \infty} \mathbf{MP}(\pi[k])$. A weighted graph equipped with these payoffs is called a *total-payoff game* or a *mean-payoff game*, respectively.

A *strategy* for Max (respectively, Min) in a game $\mathcal{G} = \langle V, E, \omega, \mathbf{P} \rangle$ (with $\mathbf{P}$ one of the previous payoffs), is a mapping $\sigma \colon V^* V_{\mathsf{Max}} \to V$ (respectively, $\sigma \colon V^* V_{\mathsf{Min}} \to V$) such that for all sequences $\pi = v_0 \cdots v_k$ with $v_k \in V_{\mathsf{Max}}$ (respectively, $v_k \in V_{\mathsf{Min}}$), $(v_k, \sigma(\pi)) \in E$. A play or finite play $\pi = v_0 v_1 \cdots$ conforms to a strategy σ of Max (respectively, Min) if for all k such that $v_k \in V_{\mathsf{Max}}$ (respectively, $v_k \in V_{\mathsf{Min}}$), $v_{k+1} = \sigma(\pi[k])$. A strategy σ is *memoryless*

[1] An example of practical application would be to perform controller synthesis taking into account energy consumption. On the other hand, the problem of computing the values in certain classes of priced timed games has recently been reduced to computing the values in min-cost reachability games [3].

[2] Our results can easily be extended by substituting a lim sup for the lim inf. The lim inf is more natural since we adopt the point of view of the maximiser Max, hence the lim inf is the *worst* partial sum seen infinitely often.

if for all finite plays π, π', we have $\sigma(\pi v) = \sigma(\pi' v)$ for all v. A strategy σ is said to be *finite-memory* if it can be encoded in a deterministic Moore machine, $\langle M, m_0, \mathsf{up}, \mathsf{dec}\rangle$, where M is a finite set representing the memory of the strategy, with an initial memory content $m_0 \in M$, $\mathsf{up}\colon M \times V \to M$ is a memory-update function, and $\mathsf{dec}\colon M \times V \to V$ a decision function such that for every finite play π and vertex v, $\sigma(\pi v) = \mathsf{dec}(\mathsf{mem}(\pi v), v)$ where $\mathsf{mem}(\pi)$ is defined by induction on the length of the finite play π as follows: $\mathsf{mem}(v_0) = m_0$, and $\mathsf{mem}(\pi v) = \mathsf{up}(\mathsf{mem}(\pi), v)$. We say that $|M|$ is the *size* of the strategy.

For all strategies σ_{Max} and σ_{Min}, for all vertices v, we let $\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})$ be the outcome of σ_{Max} and σ_{Min}, defined as the unique play conforming to σ_{Max} and σ_{Min} and starting in v. Naturally, the objective of Max is to maximise its payoff. In this model of zero-sum game, Min then wants to minimise the payoff of Max. Formally, we let $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}})$ and $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Min}})$ be the respective values of the strategies, defined as (recall that $\mathbf{P}$ is either $\mathbf{TP}$ or $\mathbf{MP}$): $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}}) = \inf_{\sigma_{\mathsf{Min}}} \mathbf{P}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}}))$ and $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Min}}) = \sup_{\sigma_{\mathsf{Max}}} \mathbf{P}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}}))$. Finally, for all vertices v, we let $\underline{\mathsf{Val}}_{\mathcal{G}}(v) = \sup_{\sigma_{\mathsf{Max}}} \mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}})$ and $\overline{\mathsf{Val}}_{\mathcal{G}}(v) = \inf_{\sigma_{\mathsf{Min}}} \mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Min}})$ be the *lower* and *upper values* of v respectively. We may easily show that $\underline{\mathsf{Val}}_{\mathcal{G}} \preccurlyeq \overline{\mathsf{Val}}_{\mathcal{G}}$. We say that strategies σ^*_{Max} of Max and σ^*_{Min} of Min are optimal if, for all vertices v: $\mathsf{Val}_{\mathcal{G}}(v, \sigma^*_{\mathsf{Max}}) = \underline{\mathsf{Val}}_{\mathcal{G}}(v)$ and $\mathsf{Val}_{\mathcal{G}}(v, \sigma^*_{\mathsf{Min}}) = \overline{\mathsf{Val}}_{\mathcal{G}}(v)$ respectively. We say that a game $\mathcal{G}$ is *determined* if for all vertices v, its lower and upper values are equal. In that case, we write $\mathsf{Val}_{\mathcal{G}}(v) = \underline{\mathsf{Val}}_{\mathcal{G}}(v) = \overline{\mathsf{Val}}_{\mathcal{G}}(v)$, and refer to it as the *value* of v. If the game is clear from the context, we may drop the index $\mathcal{G}$ of all previous values. Mean-payoff and total-payoff games are known to be determined, with the existence of optimal memoryless strategies [17, 10].

Total-payoff games have been mainly considered as a refinement of mean-payoff games [10]. Indeed, if the mean-payoff value of a game is positive (respectively, negative), its total-payoff value is necessarily $+\infty$ (respectively, $-\infty$). When the mean-payoff value is 0 however, the total-payoff is necessarily different from $+\infty$ and $-\infty$, hence total-payoff games are particularly useful in this case. Deciding whether the total-payoff value of a vertex is positive can be achieved in NP $\cap$ co-NP. In [9], the complexity is refined to UP $\cap$ co-UP, and values are shown to be effectively computable solving nested fixed point equations with a strategy iteration algorithm working in exponential time in the worst case.

Our aim is to give a pseudo-polynomial algorithm solving total-payoff games. In many cases, (e.g., mean-payoff games), a successful way to obtain such an efficient algorithm is the *value iteration paradigm*. Intuitively, value iteration algorithms compute successive approximations $x_0, x_1, \ldots, x_i, \ldots$ of the game value by restricting the number of turns that the players are allowed to play: x_i is the vector of optimal values achievable when the players play at most i turns. The sequence of values is computed by means of an operator $\mathcal{F}$, letting $v_{i+1} = \mathcal{F}(v_i)$ for all i. Good properties (Scott-continuity and monotonicity) of $\mathcal{F}$ ensure convergence towards its smallest or greatest fixed point (depending on the value of x_0), which, in some cases, happens to be the value of the game. Let us briefly explain why such a simple approach fails with total-payoff games. In our case, the operator $\mathcal{F}$ is such that $\mathcal{F}(x)(v) = \max_{v' \in E(v)} \omega(v, v') + x(v')$ for all $v \in V_{\mathsf{Max}}$ and $\mathcal{F}(x)(v) = \min_{v' \in E(v)} \omega(v, v') + x(v')$ for all $v \in V_{\mathsf{Min}}$. This definition matches the intuition that x_i are optimal values after i turns.

Then, consider the example of Figure 1(a), limited to vertices $\{v_3, v_4, v_5\}$ for simplicity. Observe that there are two simple cycles with weight 0, hence the total-payoff value of this game is finite. Max has the choice between cycling into one of these two cycles. It is easy to check that Max's optimal choice is to enforce the cycle between v_4 and v_5, securing a payoff of -1 from v_4 (because of the lim inf definition of $\mathbf{TP}$). Hence, the values of x_3, x_4 and x_5 are respectively 1, -1 and 0. In this game, we have $\mathcal{F}(x_3, x_4, x_5) =$

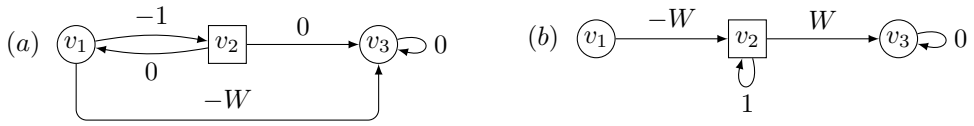

Figure 2 Two weighted graphs

$\big(2+x_4, \max(-2+x_3, -1+x_5), 1+x_4\big)$, and the vector $(1,-1,0)$ is indeed a fixed point of $\mathcal{F}$. However, it is neither the greatest nor the smallest fixed point of $\mathcal{F}$, since *if* x is a fixed point of $\mathcal{F}$, *then* $x+(a,a,a)$ is also a fixed point, for all constant $a \in \mathbb{Z}$. If we try to initialise the value iteration algorithm with value $(0,0,0)$, which could seem a reasonable choice, the sequence of computed vectors is: $(0,0,0)$, $(2,-1,1)$, $(1,0,0)$, $(2,-1,1)$, $(1,0,0)$, ... that is not stationary, and does not even contain $(1,-1,0)$. Thus, it seems difficult to compute the actual game values with an iterative algorithm relying on the $\mathcal{F}$ operator, as in the case of mean-payoff games.[3] Notice that, in the previous example, the Zwick and Paterson's algorithm [17] to solve mean-payoff games would easily conclude from the sequence above, since the vectors of interest are then the one divided by the length of the current sequence, i.e., $(0,0,0)$, $(1,-0.5,0.5)$, $(0.33,0,0)$, $(0.5,-0.25,0.25)$, $(0.2,0,0)$, ... indeed converging towards $(0,0,0)$, the mean-payoff values of this game.

Instead, as explained in the introduction, we propose a different approach that consists in reducing total-payoff games to min-cost reachability games where Min must enforce a reachability objective on top of his optimisation objective. The aim of the next section is to study these games, and we reduce total-payoff games to them in Section 4.

3 Min-cost reachability games

In this section, we consider *min-cost reachability games* (MCR games for short), a variant of total-payoff games where one player has a reachability objective that he must fulfil first, before optimising his quantitative objective. Without loss of generality, we assign the reachability objective to player Min, as this will make our reduction from total-payoff games easier to explain. Hence, when the target is not reached along a path, its payoff shall be the worst possible for Min, i.e., $+\infty$. Formally, an MCR game is played on a weighted graph $\langle V, E, \omega\rangle$ equipped with a target set of vertices $T \subseteq V$. The payoff T-$\mathbf{MCR}(\pi)$ of a play $\pi = v_0 v_1 \dots$ is given by T-$\mathbf{MCR}(\pi) = +\infty$ if the play avoids T, i.e., if for all $k \geqslant 0$, $v_k \notin T$, and T-$\mathbf{MCR}(\pi) = \mathbf{TP}(\pi[k])$ if k is the least position in π such that $v_k \in T$. Lower and upper values are then defined as in Section 2. By an indirect consequence of Martin's theorem [12], we can show that MCR games are also determined. Optimal strategies may however not exist, as we will see later.

As an example, consider the MCR game played on the weighted graph of Figure 2(a), where W is a positive integer and v_3 is the target. We claim that the values of vertices v_1 and v_2 are both $-W$. Indeed, consider the following strategy for Min: during each of the first W visits to v_2 (if any), go to v_1; else, go to v_3. Clearly, this strategy ensures that the target will eventually be reached, and that either (i) edge (v_1, v_3) (with weight $-W$) will eventually be traversed; or (ii) edge (v_1, v_2) (with weight -1) will be traversed at least W times. Hence,

[3] In the context of stochastic models like Markov decision processes, Strauch [14] already noticed that in the presence of arbitrary weights, the value iteration algorithm does not necessarily converge towards the accurate value: see [13, Ex. 7.3.3] for a detailed explanation.

in all plays following this strategy, the payoff will be at most $-W$. This strategy allows Min to secure $-W$, but he cannot ensure a lower payoff, since Max always has the opportunity to take the edge (v_1, v_3) (with weight $-W$) instead of cycling between v_1 and v_2. Hence, Max's optimal choice is to follow the edge (v_1, v_3) as soon as v_1 is reached, securing a payoff of $-W$. The Min strategy we have just given is optimal, and there is *no optimal memoryless strategy* for Min. Indeed, always playing (v_2, v_3) does not ensure a payoff $\leqslant -W$; and, always playing (v_2, v_1) does not guarantee to reach the target, and this strategy has thus value $+\infty$.

Let us note that Björklund and Vorobyov introduce in [1] the *longest shortest path problem* (LSP for short) and propose a pseudo-polynomial time algorithm to solve it. However, their definition has several subtle but important differences to ours, such as definition of the payoff of a play (equivalently, the length of a path). As an example, in the game of Figure 2(a), the play $\pi = (v_1 v_2)^\omega$ (that never reaches the target) has length $-\infty$ in their setting, while, in our setting, $\{v_3\}$-$\mathbf{MCR}(\pi) = +\infty$. Moreover, even if a pre-treatment would hypothetically allow one to use the LSP algorithm to solve MCR games, our solution is simpler to implement with the same worst-case complexity and heuristics only applicable to our value iteration solution. We now present our contributions for MCR games:

▶ **Theorem 1.** *Let $\mathcal{G} = \langle V, E, \omega, T\text{-}\mathbf{MCR}\rangle$ be an MCR game.*

1. *For $v \in V$, deciding whether $\mathsf{Val}(v) = +\infty$ can be done in polynomial time.*
2. *For $v \in V$, deciding whether $\mathsf{Val}(v) = -\infty$ is as hard as mean-payoff, in $\mathrm{NP} \cap \text{co-NP}$ and can be achieved in pseudo-polynomial time.*
3. *If $\mathsf{Val}(v) \neq -\infty$ for all vertices $v \in V$, then both players have optimal strategies. Moreover, Max always has a memoryless optimal strategy, while Min may require finite (pseudo-polynomial) memory in his optimal strategy.*
4. *Computing all values $\mathsf{Val}(v)$ (for $v \in V$), as well as optimal strategies (if they exist) for both players, can be done in (pseudo-polynomial) time $O(|V|^2|E|W)$.*

To prove the first item it suffices to notice that vertices with value $+\infty$ are exactly those from which Min cannot reach the target. Therefore the problem reduces to deciding the winner in a classical reachability game, that can be solved in polynomial time [16], using the classical *attractor* construction: in vertices of value $+\infty$, Min may play indifferently, while Max has an optimal memoryless strategy consisting in avoiding the attractor.

To prove the second item, it suffices first to notice that vertices with value $-\infty$ are exactly those with a value < 0 in the mean-payoff game played on the same graph. On the other hand, we can show that any mean-payoff game can be transformed (in polynomial time) into an MCR game such that a vertex has value < 0 in the mean-payoff game if and only if the value of its corresponding vertex in the MCR game is $-\infty$. The rest of this section focuses on the proof of the third and fourth items. We start by explaining how to compute the values in pseudo-polynomial, and we discuss optimal strategies afterward.

Computing the values. From now on, we assume, without loss of generality, that there is exactly one target vertex denoted by $\mathtt{t}$, and the only outgoing edge from $\mathtt{t}$ is a self loop with weight 0: this is reflected by denoting $\mathbf{MCR}$ the payoff mapping $\{\mathtt{t}\}$-$\mathbf{MCR}$. Our value iteration algorithm for MCR games is given in Algorithm 1. To establish its correctness, we rely mainly on the operator $\mathcal{F}$, which denotes the function $\mathbb{Z}_\infty^V \to \mathbb{Z}_\infty^V$ mapping every vector $x \in \mathbb{Z}_\infty^V$ to $\mathcal{F}(x)$ defined by $\mathcal{F}(x)(\mathtt{t}) = 0$ and

$$\mathcal{F}(x)(v) = \begin{cases} \max\limits_{v' \in E(v)} \big(\omega(v, v') + x(v')\big) & \text{if } v \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\} \\ \min\limits_{v' \in E(v)} \big(\omega(v, v') + x(v')\big) & \text{if } v \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\} \end{cases}$$

Algorithm 1: Value iteration for min-cost reachability games

Input: MCR game $\langle V, E, \omega, \mathbf{MCR}\rangle$, W largest weight in absolute value

1 $\mathsf{X}(\mathtt{t}) := 0$
2 **foreach** $v \in V \setminus \{\mathtt{t}\}$ **do** $\mathsf{X}(v) := +\infty$
3 **repeat**
4 $\quad \mathsf{X}_{pre} := \mathsf{X}$
5 $\quad$ **foreach** $v \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\}$ **do** $\mathsf{X}(v) := \max_{v' \in E(v)} \big(\omega(v,v') + \mathsf{X}_{pre}(v')\big)$
6 $\quad$ **foreach** $v \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\}$ **do** $\mathsf{X}(v) := \min_{v' \in E(v)} \big(\omega(v,v') + \mathsf{X}_{pre}(v')\big)$
7 $\quad$ **foreach** $v \in V \setminus \{\mathtt{t}\}$ such that $\mathsf{X}(v) < -(|V|-1)W$ **do** $\mathsf{X}(v) := -\infty$
8 **until** $\mathsf{X} = \mathsf{X}_{pre}$
9 **return** X

More precisely, we are interested in the sequence of iterates $x_i = \mathcal{F}(x_{i-1})$ of $\mathcal{F}$ from the initial vector x_0 defined by $x_0(v) = +\infty$ for all $v \neq \mathtt{t}$, and $x_0(\mathtt{t}) = 0$. The intuition behind the sequence $(x_i)_{i\geqslant 0}$ is that x_i *is the value of the game if we impose that* Min *must reach the target within* i *steps* (and get a payoff of $+\infty$ if he fails to do so). Formally, for a play $\pi = v_0 v_1 \cdots v_i \cdots$, we let $\mathbf{MCR}^{\leqslant i}(\pi) = \mathbf{MCR}(\pi)$ if $v_k = \mathtt{t}$ for some $k \leqslant i$, and $\mathbf{MCR}^{\leqslant i}(\pi) = +\infty$ otherwise. We further let $\overline{\mathsf{Val}}^{\leqslant i}(v) = \inf_{\sigma_{\mathsf{Min}}} \sup_{\sigma_{\mathsf{Max}}} \mathbf{MCR}^{\leqslant i}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}}))$ (where σ_{Max} and σ_{Min} are respectively strategies of Max and Min). We can show that the operator $\mathcal{F}$ allows one to compute the sequence $(\overline{\mathsf{Val}}^{\leqslant i})_{i\geqslant 0}$, i.e., for all $i \geqslant 0$: $x_i = \overline{\mathsf{Val}}^{\leqslant i}$.

Let us first show that the algorithm is *correct* when the values of all nodes are *finite*. Thanks to this characterisation, and by definition of $\overline{\mathsf{Val}}^{\leqslant i}$, it is easy to see that, for all $i \geqslant 0$: $x_i = \overline{\mathsf{Val}}^{\leqslant i} \succcurlyeq \overline{\mathsf{Val}} = \mathsf{Val}$. Moreover, $\mathcal{F}$ is a monotonic operator over the complete lattice $\mathbb{Z}_\infty^V$. By Knaster-Tarski's theorem, the fixed points of $\mathcal{F}$ form a complete lattice and $\mathcal{F}$ admits a greatest fixed point. By Kleene's fixed point theorem, using the Scott-continuity of $\mathcal{F}$, this greatest fixed point can be obtained as the limit of the non-increasing sequence of iterates $(\mathcal{F}^i(\overline{x}))_{i\geqslant 0}$ starting in the maximal vector $\overline{x}$ defined by $\overline{x}(v) = +\infty$ for all $v \in V$. As $x_0 = \mathcal{F}(\overline{x})$, the sequence $(x_i)_{i\geqslant 0}$ is also non-increasing (i.e., $x_i \succcurlyeq x_{i+1}$, for all $i \geqslant 0$) and converges towards the greatest fixed point of $\mathcal{F}$. We can further show that the value of the game Val is actually the greatest fixed point of $\mathcal{F}$. Moreover, we can bound the number of steps needed to reach that fixed point (when all values are finite—this is the point where this hypothesis is crucial), by carefully observing the possible vectors that can be computed by the algorithm: the sequence $(x_i)_{i\geqslant 0}$ is non-increasing, and stabilises after at most $(2|V|-1)W|V| + |V|$ steps on Val.

Thus, computing the sequence $(x_i)_{i\geqslant 0}$ up to stabilisation yields the values of all vertices in an MCR game *if all values are finite*. Were it not for line 7, Algorithm 1 would compute exactly this sequence. We claim that Algorithm 1 is correct even when vertices have values in $\{-\infty, +\infty\}$. Line 7 allows to cope with vertices whose value is $-\infty$: when the algorithm detects that Min can secure a value small enough from a vertex v, it sets v's value to $-\infty$. Intuitively, this is correct because if Min can guarantee a payoff smaller than $-(|V|-1) \times W$, he can force a negative cycle from which he can reach $\mathtt{t}$ with an arbitrarily small value. Hence, one can ensure that, after i iterations of the loop, $x_{i-1} \succcurlyeq \mathsf{X} \succcurlyeq \mathsf{Val}$, and the sequence still converges to Val, the greatest fixed point of $\mathcal{F}$. Finally, if some vertex v has value $+\infty$, one can check that $\mathsf{X}(v) = +\infty$ is an invariant of the loop. From that point, one can prove the correctness of the algorithm. Thus, the algorithm executes $O(|V|^2 W)$ iterations. Since each iteration can be performed in $O(|E|)$, the algorithm has a complexity of $O(|V|^2|E|W)$,

as announced in Theorem 1. As an example, consider the min-cost reachability game of Figure 2(a). The successive values for vertices (v_1, v_2) (value of the target v_3 is always 0) computed by the value iteration algorithm are the following: $(+\infty, +\infty)$, $(+\infty, 0)$, $(-1, 0)$, $(-1, -1)$, $(-2, -1)$, $(-2, -2), \ldots, (-W, -W+1)$, $(-W, -W)$. This requires $2W$ steps to converge (hence a pseudo-polynomial time).

Computing optimal strategies for both players. We now turn to the proof of the third item of Theorem 1, supposing that every vertex v of the game has a finite value $\mathsf{Val}(v) \in \mathbb{Z}$ (the case where $\mathsf{Val}(v) = +\infty$ is delt with the attractor construction).

Observe first that, Min may *need memory* to play optimally, as already shown by the example in Figure 2(a), where the target is v_3. Nevertheless, let us briefly explain why optimal strategies for Min always exist, with a memory of pseudo-polynomial size. We extract from the sequence $(x_i)_{i \geqslant 0}$ defined above (or equivalently, from the sequence of vectors X of Algorithm 1) the optimal strategy σ^*_{Min} as follows. Let k be the first index such that $x_{k+1} = x_k$. Then, for every play π ending in vertex $v \in V_{\mathsf{Min}}$, we let $\sigma^*_{\mathsf{Min}}(\pi) = \mathrm{argmin}_{v' \in E(v)} \big(\omega(v, v') + x_{k-|\pi|-1}(v')\big)$, if $|\pi| < k$, and $\sigma^*_{\mathsf{Min}}(\pi) = \mathrm{argmin}_{v' \in E(v)} \big(\omega(v, v') + x_0(v')\big)$ otherwise (those argmin may not be unique, but we can indifferently pick any of them). Since σ_{Min} only requires to know the last vertex and the length of the prefix up to k, and since $k \leqslant (2|V|-1)W|V|+|V|$ as explained above, σ^*_{Min} needs a memory of pseudo-polynomial size only. Moreover, it can be computed with the sequence of vectors X in Algorithm 1. It is not difficult to verify by induction that this strategy is optimal for Min. While optimal, this strategy might not be practical, for instance, in the framework of controller synthesis. Implementing it would require to store the full sequence $(x_i)_{i \geqslant 0}$ up to convergence step k (possibly pseudo-polynomial) in a table, and to query this large table each time the strategy is called. Instead, an alternative optimal strategy σ'_{Min} can be construct, that consists in playing successively two *memoryless strategies* σ^1_{Min} and σ^2_{Min} (σ^2_{Min} being given by the attractor construction). To determine when to switch from σ^1_{Min} to σ^2_{Min}, σ'_{Min} maintains a counter that is stored in a *polynomial* number of bits, thus the memory footprints of σ'_{Min} and σ^*_{Min} are comparable. However, σ'_{Min} is easier to implement, because σ^1_{Min} and σ^2_{Min} can be described by a pair of tables of linear size, and, apart from querying those tables, σ'_{Min} consists only in incrementing and testing the counter to determine when to switch. Moreover, this succession of two memoryless strategies allows us to also get some interesting strategy in case of vertices with values $-\infty$: indeed, we can still compute this pair of strategies, and simply modify the switching policy to run for a sufficiently long time to guarantee a value less than a given threshold. In the following, we call such a strategy a *switching strategy*.

Finally, we can show that, contrary to Min, Max always has a *memoryless optimal strategy* σ^*_{Max} defined by $\sigma^*_{\mathsf{Max}}(\pi) = \mathrm{argmax}_{v' \in E(v)} \big(\omega(v, v') + \mathsf{Val}(v')\big)$ for all finite plays π ending in $v \in V_{\mathsf{Max}}$. For example, in the game of Figure 2(a), $\sigma^*_{\mathsf{Max}}(\pi v_2) = v_3$ for all π, since $\mathsf{Val}(v_3) = 0$ and $\mathsf{Val}(v_1) = -W$. Moreover, the previously described optimal strategies can be computed along the execution of Algorithm 1. Finally, we can show that, for all vertices v, the pair of optimal strategies we have just defined yields a play $\mathsf{Play}(v, \sigma^*_{\mathsf{Max}}, \sigma^*_{\mathsf{Min}})$ which is *non-looping*, i.e., never visits the same vertex twice before reaching the target. For instance, still in the game of Figure 2(a), $\mathsf{Play}(v_1, \sigma^*_{\mathsf{Max}}, \sigma^*_{\mathsf{Min}}) = v_1 v_2 v_3^\omega$.

4 An efficient algorithm to solve total-payoff games

We now turn our attention back to total-payoff games (without reachability objective), and discuss our main contribution. Building on the results of the previous section, we introduce

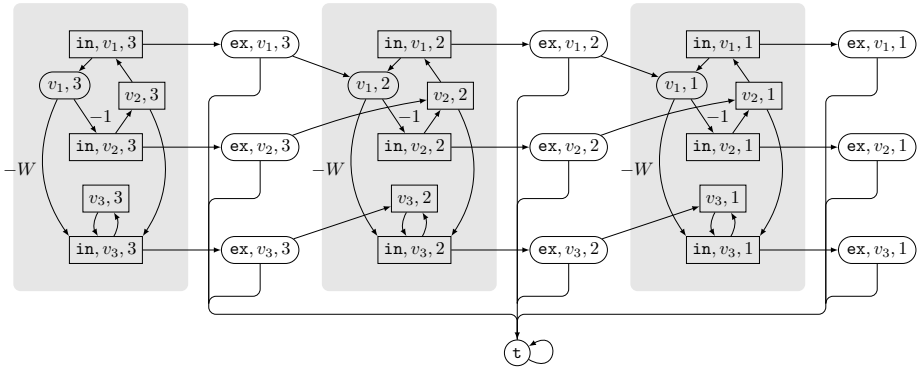

Figure 3 MCR game $\mathcal{G}^3$ associated with the total-payoff game of Figure 2(a)

the *first* (as far as we know) *pseudo-polynomial time algorithm* for solving those games in the presence of arbitrary weights, thanks to a reduction from total-payoff games to min-cost reachability games. The MCR game produced by the reduction has size pseudo-polynomial in the size of the original total-payoff game. Then, we show how to compute the values of the total-payoff game without building the entire MCR game, and explain how to deduce memoryless optimal strategies from the computation of our algorithm.

Reduction to min-cost reachability games. We provide a transformation from a total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP}\rangle$ to a min-cost reachability game $\mathcal{G}^K$ such that the values of $\mathcal{G}$ can be extracted from the values in $\mathcal{G}^K$ (as formalised below). Intuitively, $\mathcal{G}^K$ simulates the game where players play in $\mathcal{G}$; Min may propose to stop playing and reach a fresh vertex $\mathtt{t}$ acting as the target; Max can then accept, in which case we reach the target, or refuse at most K times, in which case the game continues. Structurally, $\mathcal{G}^K$ consists of a sequence of copies of $\mathcal{G}$ along with some new states that we now describe formally. We let $\mathtt{t}$ be a fresh vertex, and, for all $n \geqslant 1$, we define the min-cost reachability game $\mathcal{G}^n = \langle V^n, E^n, \omega^n, \{\mathtt{t}\}\text{-}\mathbf{MCR}\rangle$ where V^n_{Max} (respectively, V^n_{Min}) consists of n copies (v, j), with $1 \leqslant j \leqslant n$, of each vertex $v \in V_{\mathsf{Max}}$ (respectively, $v \in V_{\mathsf{Min}}$) and some *exterior vertices* $(\mathtt{ex}, v, j)$ for all $v \in V$ and $1 \leqslant j \leqslant n$ (respectively, *interior vertices* $(\mathtt{in}, v, j)$ for all $v \in V$ and $1 \leqslant j \leqslant n$). Moreover, V^n_{Max} contains the fresh target vertex $\mathtt{t}$. Edges are given by

$$\begin{aligned} E^n = {} & \{(\mathtt{t},\mathtt{t})\} \uplus \big\{\big((v,j),(\mathtt{in},v',j)\big) \mid (v,v') \in E, 1 \leqslant j \leqslant n\big\} \\ & \uplus \big\{\big((\mathtt{in},v,j),(v,j)\big) \mid v \in V, 1 \leqslant j \leqslant n\big\} \uplus \big\{\big((\mathtt{ex},v,j),\mathtt{t}\big) \mid v \in V, 1 \leqslant j \leqslant n\big\} \\ & \uplus \big\{\big((\mathtt{in},v,j),(\mathtt{ex},v,j)\big) \mid v \in V, 1 \leqslant j \leqslant n\big\} \\ & \uplus \big\{\big((\mathtt{ex},v,j),(v,j-1)\big) \mid v \in V, 1 < j \leqslant n\big\}. \end{aligned}$$

All edge weights are zero, except edges $\big((v,j),(\mathtt{in},v',j)\big)$ that have weight $\omega(v,v')$.

For example, considering the weighted graph of Figure 2(a), the corresponding reachability total-payoff game $\mathcal{G}^3$ is depicted in Figure 3 (where weights 0 have been removed). The next proposition formalises the relationship between the two games.

▶ Proposition 2. Let $K = |V|(2(|V|-1)W+1)$. For all $v \in V$ and $k \geqslant K$,

- $\mathsf{Val}_{\mathcal{G}}(v) \neq +\infty$ if and only if $\mathsf{Val}_{\mathcal{G}}(v) = \mathsf{Val}_{\mathcal{G}^k}((v,k))$;
- $\mathsf{Val}_{\mathcal{G}}(v) = +\infty$ if and only if $\mathsf{Val}_{\mathcal{G}^k}((v,k)) \geqslant (|V|-1)W+1$.

The bound K is found by using the fact (informally described in the previous section) that if not infinite, the value of a min-cost reachability game belongs in $[-(|V|-1)\times W+1, |V|\times W]$,

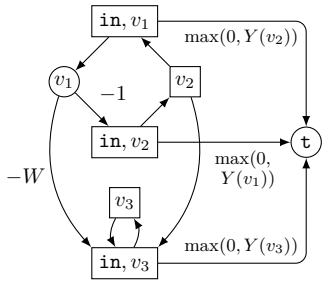

Figure 4 MCR game $\mathcal{G}_Y$ associated with the total-payoff game of Figure 2(a)

and that after enough visits of the same vertex, an adequate loop ensures that $\mathcal{G}^k$ verifies the above properties.

Value iteration algorithm for total-payoff games. By Proposition 2, an immediate way to obtain a value iteration algorithm for total-payoff games is to build game $\mathcal{G}^K$, run Algorithm 1 on it, and map the computed values back to $\mathcal{G}$. We take advantage of the structure of $\mathcal{G}^K$ to provide a better algorithm that avoids building $\mathcal{G}^K$. We first compute the values of the vertices in the *last copy of the game* (vertices of the form $(v, 1)$, $(\texttt{in}, v, 1)$ and $(\texttt{ex}, v, 1)$), then of those in the penultimate (vertices of the form $(v, 2)$, $(\texttt{in}, v, 2)$ and $(\texttt{ex}, v, 2)$), and so on.

We formalise this idea as follows. Let Z^j be a vector mapping each vertex v of $\mathcal{G}$ to the value $Z^j(v)$ of vertex (v, j) in $\mathcal{G}^K$. Then, let us define an operator $\mathcal{H}$ such that $Z^{j+1} = \mathcal{H}(Z^j)$. The intuition behind the definition of $\mathcal{H}(Y)$ for some vector Y, is to extract from $\mathcal{G}^K$ one copy of the game, and make Y appear in the weights of some edges as illustrated in Figure 4. This game, $\mathcal{G}_Y$, simulates a play in $\mathcal{G}$ in which Min can opt for 'leaving the game' at each round (by moving to the target), obtaining $\max(0, Y(v))$, if v is the current vertex. Then $\mathcal{H}(Y)(v)$ is defined as the value of v in $\mathcal{G}_Y$. By construction, it is easy to see that $Z^{j+1} = \mathcal{H}(Z^j)$ holds for all $j \geqslant 1$. Furthermore, we define $Z^0(v) = -\infty$ for all v, and have $Z^1 = \mathcal{H}(Z^0)$. One can prove the following properties of $\mathcal{H}$: (i) $\mathcal{H}$ is monotonic, but may not be Scott-continuous; (ii) the sequence $(Z^j)_{j \geqslant 0}$ converges towards $\mathsf{Val}_{\mathcal{G}}$.

We are now ready to introduce Algorithm 2 to solve total-payoff games. Intuitively, the outer loop computes, in variable Y, a non-decreasing sequence of vectors whose limit is $\mathsf{Val}_{\mathcal{G}}$, and that is stationary (this is not necessarily the case for the sequence $(Z^j)_{j \geqslant 0}$). Line 1 initialises Y to Z^0. Each iteration of the outer loop amounts to running Algorithm 1 to compute $\mathcal{H}(\mathsf{Y}_{pre})$ (lines 3 to 10), then detecting if some vertices have value $+\infty$, updating Y accordingly (line 11, following the second item of Proposition 2). One can show that, for all $j > 0$, if we let Y^j be the value of Y after the j-th iteration of the main loop, $Z^j \preccurlyeq Y^j \preccurlyeq \mathsf{Val}_{\mathcal{G}}$, which ensures the correctness of the algorithm.

▶ **Theorem 3.** *If a total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP} \rangle$ is given as input, Algorithm 2 outputs the vector $\mathsf{Val}_{\mathcal{G}}$ of optimal values, after at most $K = |V|(2(|V| - 1)W + 1)$ iterations of the external loop. The complexity of the algorithm is $O(|V|^4|E|W^2)$.*

The number of iterations in each internal loop is controlled by Theorem 1. On the example of Figure 2(a), only 2 external iterations are necessary, but the number of iterations of each internal loop would be $2W$. By contrast, for the total-payoff game depicted in Figure 2(b), each internal loop requires 2 iterations to converge, but the external loop takes W iterations to stabilise. A combination of both examples would experience a pseudo-polynomial number of iterations to converge in both the internal and external loops, matching the W^2 term of the above complexity.

Algorithm 2: A value iteration algorithm for total-payoff games

Input: Total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP}\rangle$, W largest weight in absolute value

foreach $v \in V$ **do** $\mathsf{Y}(v) := -\infty$
repeat
  foreach $v \in V$ **do** $\mathsf{Y}_{pre}(v) := \mathsf{Y}(v)$; $\mathsf{Y}(v) := \max(0, \mathsf{Y}(v))$; $\mathsf{X}(v) := +\infty$
  repeat
    $\mathsf{X}_{pre} := \mathsf{X}$
    foreach $v \in V_{\mathsf{Max}}$ **do** $\mathsf{X}(v) := \max_{v' \in E(v)} \left[\omega(v, v') + \min(\mathsf{X}_{pre}(v'), \mathsf{Y}(v'))\right]$
    foreach $v \in V_{\mathsf{Min}}$ **do** $\mathsf{X}(v) := \min_{v' \in E(v)} \left[\omega(v, v') + \min(\mathsf{X}_{pre}(v'), \mathsf{Y}(v'))\right]$
    foreach $v \in V$ such that $\mathsf{X}(v) < -(|V| - 1)W$ **do** $\mathsf{X}(v) := -\infty$
  until $\mathsf{X} = \mathsf{X}_{pre}$
  $\mathsf{Y} := \mathsf{X}$
  foreach $v \in V$ such that $\mathsf{Y}(v) > (|V| - 1)W$ **do** $\mathsf{Y}(v) := +\infty$
until $\mathsf{Y} = \mathsf{Y}_{pre}$
return Y

Optimal strategies. In Section 3, we have shown, for any min-cost reachability game, the existence of a pair of memoryless strategies permitting to reconstruct a *switching* optimal strategy for Min (if every vertex has value different from $-\infty$, or a strategy ensuring any possible threshold for vertices with value $-\infty$). If we apply this construction to the game $\mathcal{G}_{\mathsf{Val}_\mathcal{G}}$, we obtain a pair $(\sigma^1_{\mathsf{Min}}, \sigma^2_{\mathsf{Min}})$ of strategies (remember that σ^2_{Min} is a strategy obtained by the attractor construction, hence it will not be useful for us for total-payoff games). Consider the strategy $\overline{\sigma_{\mathsf{Min}}}$, obtained by projecting σ^1_{Min} on V as follows: for all finite plays π and vertex $v \in V_{\mathsf{Min}}$, let $\overline{\sigma_{\mathsf{Min}}}(\pi v) = v'$ if $\sigma^1_{\mathsf{Min}}(v) = (\mathtt{in}, v')$. We can show that $\overline{\sigma_{\mathsf{Min}}}$ is optimal for Min in $\mathcal{G}$. Notice that σ^1_{Min}, and hence $\overline{\sigma_{\mathsf{Min}}}$, can be computed during the last iteration of the value iteration algorithm, as explained in the case of min-cost reachability. A similar construction can be done to compute an optimal strategy of Max.

5 Implementation and heuristics

In this section, we report on a prototype implementation of our algorithms.[4] For convenience reasons, we have implemented them as an add-on to PRISM-games [5], although we could have chosen to extend another model-checker as we do not rely on the probabilistic features of PRISM models (i.e., we use the PRISM syntax of *stochastic multi-player games*, allowing arbitrary rewards, and forbidding probability distributions different of Dirac ones). We then use rPATL specifications of the form $\langle\langle C\rangle\rangle \mathsf{R}^{\min\,/\,\max=?}[\mathsf{F}^\infty \varphi]$ and $\langle\langle C\rangle\rangle \mathsf{R}^{\min\,/\,\max=?}[\mathsf{F}^c \bot]$ to model respectively min-cost reachability games and total-payoff games, where C represents a coalition of players that want to minimise/maximise the payoff, and φ is another rPATL formula describing the target set of vertices (for total-payoff games, such a formula is not necessary). We have tested our implementation on toy examples. On the parametric one studied after Theorem 3, obtained by mixing the graphs of Figure 2 and repeating them for n layers, results obtained by applying our algorithm for total-payoff games are summarised in Table 1, where for each pair (W, n), we give the time t in seconds, the number k_e of iterations

[4] Source and binary files, as well as some examples, can be downloaded from `http://www.ulb.ac.be/di/verif/monmege/tool/TP-MCR/`.

Table 1 Results of value iteration on a parametric example

		without heuristics			with heuristics		
W	n	t	k_e	k_i	t	k_e	k_i
50	100	0.52s	151	12,603	0.01s	402	1,404
50	500	9.83s	551	53,003	0.42s	2,002	7,004
200	100	2.96s	301	80,103	0.02s	402	1,404
200	500	45.64s	701	240,503	0.47s	2,002	7,004
500	1,000	536s	1,501	1,251,003	2.37s	4,002	14,004

in the external loop, and the total number k_i of iterations in the internal loop.

We close this section by sketching two techniques that can be used to speed up the computation of the fixed point in Algorithms 1 and 2. We fix a weighted graph $\langle V, E, \omega \rangle$. Both accelerations rely on a topological order of the strongly connected components (SCC for short) of the graph, given as a function $\mathsf{c}\colon V \to \mathbb{N}$, mapping each vertex to its *component*, verifying that (i) $\mathsf{c}(V) = \{0, \dots, p\}$ for some $p \geqslant 0$, (ii) $\mathsf{c}^{-1}(q)$ is a maximal SCC for all q, (iii) and $\mathsf{c}(v) \geqslant \mathsf{c}(v')$ for all $(v, v') \in E$.[5] In case of an MRC game with $\mathtt{t}$ the unique target, $\mathsf{c}^{-1}(0) = \{\mathtt{t}\}$. Intuitively, c induces a directed acyclic graph whose vertices are the sets $\mathsf{c}^{-1}(q)$ for all $q \in \mathsf{c}(V)$, and with an edge (S_1, S_2) if and only if there are $v_1 \in S_1, v_2 \in S_2$ such that $(v_1, v_2) \in E$.

The *first acceleration heuristic* is a divide-and-conquer technique that consists in applying Algorithm 1 (or the inner loop of Algorithm 2) iteratively on each $\mathsf{c}^{-1}(q)$ for $q = 0, 1, 2, \dots, p$, using at each step the information computed during steps $j < q$ (since the value of a vertex v depends only on the values of the vertices v' such that $\mathsf{c}(v') \leqslant \mathsf{c}(v)$). The *second acceleration heuristic* consists in studying more precisely each component $\mathsf{c}^{-1}(q)$. Having already computed the optimal values $\mathsf{Val}(v)$ of vertices $v \in \mathsf{c}^{-1}(\{0, \dots, q-1\})$, we ask an oracle to precompute a finite set $S_v \subseteq \mathbb{Z}_\infty$ of possible optimal values for each vertex $v \in \mathsf{c}^{-1}(q)$. For MCR games and the inner iteration of the algorithm for total-payoff games, one way to construct such a set S_v is to consider that possible optimal values are the one of non-looping paths inside the component exiting it, since, in MCR games, there exist optimal strategies for both players whose outcome is a non-looping path (see Section 3).

We can identify classes of weighted graphs for which there exists an oracle that runs in polynomial time and returns, for all vertices v, a set S_v of polynomial size. On such classes, Algorithms 1 and 2, enhanced with our two acceleration techniques, *run in polynomial time*. For instance, for all fixed positive integers L, the class of weighted graphs where every component $\mathsf{c}^{-1}(q)$ uses at most L distinct weights (that can be arbitrarily large in absolute value) satisfies this criterion. Table 1 contains the results obtained with the heuristics on the parametric example presented before. Observe that the acceleration technique permits here to decrease drastically the execution time, the number of iterations in both loops depending not even anymore on W. Even though the number of iterations in the external loop increases with heuristics, due to the decomposition, less computation is required in each internal loop since we only apply the computation for the active component.

[5] Such a mapping is computable in linear time, e.g., by Tarjan's algorithm [15].

References

1 Henrik Björklund and Sergei Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics*, 155:210–229, 2007.

2 Thomas Brihaye, Gilles Geeraerts, Axel Haddad, and Benjamin Monmege. To reach or not to reach? Efficient algorithms for total-payoff games. Research Report 1407.5030, arXiv, July 2014.

3 Thomas Brihaye, Gilles Geeraerts, Shankara Narayanan Krishna, Lakshmi Manasa, Benjamin Monmege, and Ashutosh Trivedi. Adding negative prices to priced timed games. In *Proceedings of the 25th International Conference on Concurrency Theory (CONCUR'14)*, volume 8704 of *Lecture Notes in Computer Science*, pages 560–575. Springer, 2014.

4 Luboš Brim, Jakub Chaloupka, Laurent Doyen, Rafaella Gentilini, and Jean-François Raskin. Faster algorithms for mean-payoff games. *Formal Methods for System Design*, 38(2):97–118, 2011.

5 Taolue Chen, Vojtěch Forejt, Marta Kwiatkowska, David Parker, and Aistis Simaitis. Automatic verification of competitive stochastic systems. *Formal Methods in System Design*, 43(1):61–92, 2013.

6 Carlo Comin and Romeo Rizzi. An improved pseudo-polynomial upper bound for the value problem and optimal strategy synthesis in mean payoff games. Technical Report 1503.04426, arXiv, 2015.

7 Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.

8 Emmanuel Filiot, Rafaella Gentilini, and Jean-François Raskin. Quantitative languages defined by functional automata. In *Proceedings of the 23rd International Conference on Concurrency theory (CONCUR'12)*, volume 7454 of *Lecture Notes in Computer Science*, pages 132–146. Springer, 2012.

9 Thomas Martin Gawlitza and Helmut Seidl. Games through nested fixpoints. In *Proceedings of the 21st International Conference on Computer Aided Verification (CAV'09)*, volume 5643 of *Lecture Notes in Computer Science*, pages 291–305. Springer, 2009.

10 Hugo Gimbert and Wiesław Zielonka. When can you play positionally? In *Proceedings of the 29th International Conference on Mathematical Foundations of Computer Science (MFCS'04)*, volume 3153 of *Lecture Notes in Computer Science*, pages 686–698. Springer, 2004.

11 Leonid Khachiyan, Endre Boros, Konrad Borys, Khaled Elbassioni, Vladimir Gurvich, Gabor Rudolf, and Jihui Zhao. On short paths interdiction problems: Total and node-wise limited interdiction. *Theory of Computing Systems*, 43:204–233, 2008.

12 Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

13 Martin L. Puterman. *Markov Decision Processes*. John Wiley & Sons, Inc., New York, NY, 1994.

14 Ralph E. Strauch. Negative dynamic programming. *The Annals of Mathematical Statistics*, 37:871–890, 1966.

15 Robert E. Tarjan. Depth first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972.

16 Wolfgang Thomas. On the synthesis of strategies in infinite games. In *Symposium on Theoretical Aspects of Computer Science (STACS'95)*, volume 900 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 1995.

17 Uri Zwick and Michael S. Paterson. The complexity of mean payoff games. *Theoretical Computer Science*, 158:343–359, 1996.

A Determinacy of total-payoff and min-cost reachability games

Consider a quantitative game $\mathcal{G} = \langle V, E, \omega, \mathbf{P}\rangle$ and a vertex $v \in V$. We will prove the determinacy result by using the Borel determinacy result of [12]. Indeed, for an integer M, consider Win_M to be the set of infinite plays with a payoff less than or equal to M. Payoff mapping T-**MCR** is easily shown to be Borel measurable since the set of plays with finite payoff is then a countable union of open sets of plays. For **TP**, it is shown by considering a pointwise limit of Borel measurable functions. Therefore, for payoff **P** representing total-payoff and min-cost reachability, we know that Win_M is a Borel set, so that the qualitative game defined over the graph $\langle V, E, \omega\rangle$ with winning condition Win_M is determined. We now use this preliminary result to show our determinacy result.

We first consider cases where lower or upper values are infinite. Suppose first that $\underline{\mathsf{Val}}(v) = -\infty$. We have to show that $\overline{\mathsf{Val}}(v) = -\infty$ too. Let M be an integer. From $\underline{\mathsf{Val}}(v) < M$, we know that for all strategy σ_{Max} of Max, there exists a strategy σ_{Min} for Min, such that $\mathbf{P}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) \leqslant M$. In particular, Max has no winning strategy in the qualitative game equipped with Win_M as a winning condition. By the previous determinacy result, we know that Min has a winning strategy in that case, i.e., a strategy σ_{Min} such that every strategy σ_{Max} of Max verifies $\mathbf{P}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) \leqslant M$. This exactly means that $\overline{\mathsf{Val}}(v) \leqslant M$. Since this holds for every value M, we get that $\overline{\mathsf{Val}}(v) = -\infty$. The proof goes exactly in a symetrical way to show that $\overline{\mathsf{Val}}(v) = +\infty$ implies $\underline{\mathsf{Val}}(v) = +\infty$.

Consider then the case where both $\overline{\mathsf{Val}}(v)$ and $\underline{\mathsf{Val}}(v)$ are finite values. Suppose that $\underline{\mathsf{Val}}(v) < \overline{\mathsf{Val}}(v)$ and consider a real number r strictly in-between those two values. From $r < \overline{\mathsf{Val}}(v)$, we deduce that Min has no winning strategy from v in the qualitative game with winning condition Win_r. Identically, from $\underline{\mathsf{Val}}(v) < r$, we deduce that Max has no winning strategy from v in the same game. This contradicts the determinacy of this qualitative game. Hence, $\underline{\mathsf{Val}}(v) = \overline{\mathsf{Val}}(v)$.

B Comparison with the *longest shortest path problem* of Björklund and Vorobyov

Björklund and Vorobyov have studied games related to our min-cost reachability games, as the longest shortest path problem (LSP for short) in [1]. To clarify the following discussion, let us recall the definition of LSP, adapted to our syntax.

The LSP problem considers a weighted graph $\mathcal{G}$, whose vertices are partitioned into two players, and equipped with a single target vertex $\mathtt{t}$. The problem asks to find a *memoryless* strategy σ_{Max} of Max such that in the graph $\mathcal{G}_{\sigma_{\mathsf{Max}}}$, obtained from $\mathcal{G}$ by deleting all outgoing edges from vertices of V_{Max} except those selected in σ_{Max}, the length of the shortest path from every vertex to the target $\mathtt{t}$ is as large as possible (over all memoryless strategies). The definition of paths from a vertex v to the target is, however, very different from ours: such a path may indeed be a finite play, without cycle, from v to the target $\mathtt{t}$ (in which case its length is the sum of the weights of edges). However, any cycle containing v (even if it cannot reach the target) is also considered as a path to the target. The *length* of such cycle is defined as follows:

1. if the cycle has a negative weight, its length is $-\infty$;
2. else, if it has a positive weight, its length is $+\infty$;
3. else, if it has zero weight, its length is 0.

Consider, as an example, the weighted graph of Fig. 5, where, once again, Max vertices are depicted with circles. Clearly, vertices $\{v_1, v_2, v_3\}$ and $\{v_4\}$ form independant subgames. We first consider the status of v_1, v_2 and v_3. In this subgame, Max has two possible strategies: σ_1 that selects the (v_1, v_2) edge, or σ_2 that selects the (v_1, v_3) edge instead.

1. If we select σ_1, we obtain the following values for v_1, v_2 and v_3 respectively: 0, 0 and 1. Indeed, from vertex v_2, we can loop in the zero cycle in-between v_1 and v_2, which is better than the simple path to the target (having weight 3), so that the distance from v_1 and v_2 is 0. From vertex v_3, only a single path leads to the target with weight 1.
2. If we select σ_2, we obtain the following values for v_1, v_2 and v_3 respectively: 1, 3 and 1. Indeed, from vertex v_3, the distance is the length of the simple path to the target, which is better than looping in-between v_1 and v_3 with a positive weight cycle. Hence, from vertex v_1, the distance is 1 too, whereas from v_2, the distance is 3.

Now, let us turn our attention to v_4. Again, we must consider two strategies: strategy σ_3 selects the (v_4, v_4) edge (the self-loop on v_4) and strategy σ_4 selects the $(v_4, \mathtt{t})$ edge.

1. If we select σ_3, we obtain value $-\infty$ for v_4, because all plays starting in v_4 loop in this negatively priced cycle. Recall that with our definition of min-cost reachability, such a cycle that never reaches the target

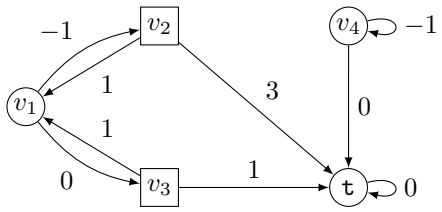

Figure 5 An instance of the longest shortest path problem

would yield a total-payoff of $+\infty$, regardless of the weights. This is coherent with our intuition that the first goal of Min is to reach the target.

2. If we select σ_4, the value of v_4 is now 0, which is thus better for Max.

Finally, the optimal strategy for Max is to choose edges (v_1, v_3) and $(v_4, \mathtt{t})$, leading to the LSP distances $(1, 3, 1, 0)$ for vertices (v_1, v_2, v_3, v_4). With our definition of min-cost reachability game, the situation is *completely different*: the value vector is then $(2, 3, 1, +\infty)$, associated with the optimal strategy σ_{Max} selecting edges (v_1, v_2) and (v_4, v_4).

Indeed, two main differences separate our two definitions. The first one is the treatment of negative weight cycles. Actually, in LSP, the fact that after selecting strategy σ_{Max} a vertex cannot reach the target anymore in $\mathcal{G}_{\sigma_{\mathsf{Max}}}$ does not prevent from mapping the distance $-\infty$ to a vertex contained in a negative cycle. This is in contrast with our definition, that would benefit to Max since in such a situation, the target is not reachable leading to the value $+\infty$.[6] The second difference consists in the treatment of zero weight cycle. In LSP, a distance zero is then computed, which is highly related with the fact that the authors want to apply the resolution of LSP to mean-payoff games.

Nevertheless, in Section 9 of [1], the authors *briefly* study another more natural definition, mapping zero weight cycle to the distance $+\infty$ (closer to our definition). In that case, the LSP distances of vertices v_1, v_2 and v_3 then match our value vector in min-cost reachability games. Unfortunately, it is not clear how the algorithm that is presented by Björklund and Vorobyov can be adapted to accomodate this new definition of the cost and compute the values of the nodes. Indeed, Proposition 9.1 of [1] proves that the decision version of the LSP (i.e., deciding if the LSP of a vertex is positive or not in a given graph) is in NP ∩ co-NP, and they claim (without formal proof) that their pseudo-polynomial time algorithm permits to decide this problem. This requires a first transformation of the problem, so that no zero weight cycles remain in the weighted graph. This transformation, explained above Proposition 9.1, is correct but *does not preserve the LSP of the vertices*, so that their algorithm do not compute the LSP distance of vertices, but only study their positivity.

Hence, the comparison between our definition of min-cost reachability and the definition of LSP can be summarised as follows:

1. Either one relies on the first definition of LSP given in the paper. In this case, Björklund and Vorobyov propose a pseudo-polynomial time algorithm to compute the LSP, but these values *do not match our definition of the min-cost reachability* payoff, as demonstrated by the above example.
2. Or one relies on the second definition. In this case, we believe that the LSP values may match our definition of min-cost reachability but the paper *does not explain formally how to compute those new values.*

It is plausible that the pseudo-polynomial time algorithm of Björklund and Vorobyov could be adapted to compute the LSP value according to the second definition. Yet, even in this case, there are several points of comparison worth mentioning:

1. The worst-case complexity of the algorithm proposed by Björklund and Vorobyov is $O(|V|^2|E|W)$, which matches ours (see Proposition 5). Nevertheless, our solution is much easier to describe and to implement: while they must make repeated calls to a modified version of the Bellman-Ford algorithm (the modification is crucial to obtain their complexity), we have a simple fixed point algorithm.
2. Our algorithm exploits the value iteration paradigm (that can be seen as a *backward induction*), while theirs is a *strategy iteration algorithm*. Because of that, there are examples on which our algorithm is more efficient that theirs (although the worst-case complexity is the same). As an example, the LSP instance presented in Fig. 2 of [1], with $2n + 1$ vertices but a biggest weight W exponential in n, requires an exponential number 2^n of iterations for the strategy iteration algorithm, but our value iteration algorithm (based on backward

[6] We believe that this difference would certainly be eliminated by our precomputation of vertices of value $+\infty$ presented in the first item of Theorem 1.

induction) would solve it in linear time with respect to n. More precisely, with our tool, we are able to compute the values of this game in less than a millisecond for constant $n = 15$ (i.e., a game with 32 vertices and largest weight $W \approx 2.68 \cdot 10^8$): the value of the initial state obtained is $4.74 \cdot 10^9$, and as aforementioned, the number of iterations that our value iteration requires is 16, linear in n. Our tool being an add-on of the PRISM model-checker, relying on the use of integer rewards, we have not been able to build the game for a value n greater than 15.

3. We propose several acceleration heuristics that perform well on several examples we have tried. These accelerations cannot easily be incorporated in their algorithm (because it is a strategy iteration algorithm).
4. Finally, from the theoretical point of view, Björklund and Vorobyov do not study the strategies of the opponent player Min. In contrast, our study allows us to produce optimal strategies for both players, and in particular, show that memoryless strategies may not be sufficient for Min, whereas they are enough for Max.

C Finding vertices of value $+\infty$ in min-cost reachability games

Let $\mathcal{G} = \langle V, E, \omega, T\text{-}\mathbf{MCR}\rangle$ be a min-cost reachability game. Notice that for all plays $\pi = v_0 v_1 \cdots$, $T\text{-}\mathbf{MCR}(\pi) = +\infty$ if and only if $v_k \notin T$ for all $k \geqslant 0$, i.e., π avoids the target. Then, let us show that the classical attractor technique [16] allows us to compute the set $V_{+\infty} = \{v \in V \mid \mathsf{Val}(v) = +\infty\}$. Recall that the attractor of a set T of vertices is obtained thanks to the sequence $\mathsf{Attr}_0(T), \ldots, \mathsf{Attr}_i(T), \ldots$ where: $\mathsf{Attr}_0(T) = T$; and for all $i \geqslant 0$:

$$\mathsf{Attr}_{i+1}(T) = \mathsf{Attr}_i(T) \cup \{v \in V_{\mathsf{Min}} \mid E(v) \cap \mathsf{Attr}_i(T) \neq \emptyset\} \cup \{v \in V_{\mathsf{Max}} \mid E(v) \subseteq \mathsf{Attr}_i(T)\}\,.$$

It is well-known that this sequence converges after at most $|V|$ steps to the set $\mathsf{Attr}(T)$ of all vertices from which Min has a memoryless strategy to ensure reaching T. Hence, under our hypothesis, $V_{+\infty} = V \setminus \mathsf{Attr}(T)$. This proves the first item of Theorem 1.

D Finding vertices of value $-\infty$ in min-cost reachability games

We give here the proof of the following aiming at computing the set $V_{-\infty}$ of vertices with a value $-\infty$ in a min-cost reachability game $\mathcal{G} = \langle V, E, \omega, T\text{-}\mathbf{MCR}\rangle$, from which we suppose that there is no more vertices with value $+\infty$.

▶ Proposition 4. For all MCR game $\mathcal{G} = \langle V, E, \omega, \mathbf{MCR}\rangle$, for all vertices v of $\mathcal{G}$, $\mathsf{Val}_{\mathcal{G}}(v) = -\infty$ if and only if $\mathsf{Val}_{\mathcal{G}'}(v) < 0$, where $\mathcal{G}'$ is the mean-payoff game $\langle V, E, \omega, \mathbf{MP}\rangle$. Conversely, given a mean-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{MP}\rangle$, we can build, in polynomial time, an MCR game $\mathcal{G}'$ such that for all vertices v of $\mathcal{G}$: $\mathsf{Val}_{\mathcal{G}}(v) < 0$ if and only if $\mathsf{Val}_{\mathcal{G}'}(v) = -\infty$.

Proof. Consider first a min-cost reachability game $\mathcal{G} = \langle V, E, \omega, T\text{-}\mathbf{MCR}\rangle$ such that $\mathsf{Val}_{\mathcal{G}}(v) \neq +\infty$ for all $v \in V$, and $\mathcal{G}' = \langle V, E, \omega, \mathbf{MP}\rangle$ the same weighted graph equipped of a mean-payoff objective.

If $\mathsf{Val}_{\mathcal{G}'}(v) < 0$, the outcome starting in v and following a profile $(\sigma^*_{\mathsf{Max}}, \sigma^*_{\mathsf{Min}})$ of optimal memoryless strategies necessarily starts with a finite prefix and then loops in a cycle with a total weight less than 0. For every $M > 0$, we construct a strategy σ_{Min} that ensures in $\mathcal{G}$ a cost less than or equal to $-M$: this will prove that $\mathsf{Val}_{\mathcal{G}}(v) = -\infty$. Since in every vertex of $\mathcal{G}'$, Min has a strategy in $\mathcal{G}$ to reach the final vertices (otherwise there would exist a vertex with value $+\infty$), there exists w'' such that Min can reach from any vertex of $\mathcal{G}'$ the target with a cost at most w''. The strategy σ_{Min} of Min is then to follow σ^*_{Min} until the accumulated cost is less than $-M - w''$ (which we can prove will happen), at which point it follows its strategy to reach the target.

Reciprocally, if $\mathsf{Val}_{\mathcal{G}}(v) = -\infty$, consider $M = |V|W$ and a strategy σ^M_{Min} of Min ensuring a cost less than $-M$, i.e., such that $\mathsf{Val}(v, \sigma^M_{\mathsf{Min}}) < -M$. Consider the finitely-branching tree built from $\mathcal{G}$ by unfolding the game from vertex v and resolving the choices of Min with strategy σ^M_{Min}. Each branch of this tree corresponds to a possible strategy of Max. Since this strategy generates a finite cost, we are certain that every such branch leads to a vertex of T. If we trim the tree at those vertices, we finally obtain a finite tree. Now, for a contradiction, suppose the optimal memoryless strategy σ^*_{Max} of Max ensures a non-negative mean-payoff, that is, $\mathsf{Val}_{\mathcal{G}'}(v, \sigma^*_{\mathsf{Max}}) \geqslant 0$. Consider the branch of the previous tree where Max follows strategy σ^*_{Max}. Since this finite branch has cost less than $-M = -|V|W$, we know for sure that there is two occurrences of the same vertex v' with an in-between weight less than 0: otherwise, by removing all nonnegative cycles, we obtain a play without repetition of vertices, henceforth of length bounded by $|V|$, and therefore of cost at least $-M$. Suppose that $v' \in V_{\mathsf{Max}}$. Then, Min has a strategy σ_{Min} to ensure a negative mean-payoff $\mathsf{Val}_{\mathcal{G}'}(v, \sigma_{\mathsf{Min}}) < 0$: indeed, it simply modifies its strategy so that it always follows his choices made in the negative cycle starting in v', ensuring that, against the optimal strategy σ^*_{Max} of Max, he gets a mean-payoff being the cost of the cycle. This is a contradiction since Max is supposed to have a strategy ensuring a non-negative mean-payoff from v. Hence, $v' \in V_{\mathsf{Min}}$. But the same contradiction

appears in that case since Min can force that it always stays in the negative cycle by modifying his strategy. Finally, we have proved that Max cannot have a memoryless strategy ensuring a non-negative mean-payoff from v. By memoryless determinacy of the mean-payoff games, this ensures that Min has a memoryless strategy ensuring a negative mean-payoff from v.

Hence, we have shown that $\mathsf{Val}_{\mathcal{G}}(v) = -\infty$ if and only if $\mathsf{Val}_{\mathcal{G}'}(v) < 0$, which concludes the first claim of Proposition 4.

Conversely, we reduce mean-payoff games to min-cost reachability games as follows. Let $\mathcal{G} = \langle V, E, \omega, \mathbf{MP}\rangle$ be a mean-payoff game. Without loss of generality, we may suppose that the graph of the game is bipartite, in the sense that $E \subseteq V_{\mathsf{Max}} \times V_{\mathsf{Min}} \cup V_{\mathsf{Min}} \times V_{\mathsf{Max}}$. The problem we are interested in is to decide whether $\mathsf{Val}_{\mathcal{G}}(v) < 0$ for a given vertex v. We now construct a min-cost reachability game $\mathcal{G}' = \langle V', E', \omega', T'\text{-}\mathbf{MCR}\rangle$ from $\mathcal{G}$. The only difference is the presence of a fresh target vertex $\mathtt{t}$ on top of vertices of V: $V' = V \uplus \{\mathtt{t}\}$ with $T' = \{\mathtt{t}\}$. Edges of $\mathcal{A}'$ are given by $E' = E \cup \{(v, \mathtt{t}) \mid v \in V_{\mathsf{Min}}\} \cup \{(\mathtt{t}, \mathtt{t})\}$. Weights of edges are given by: $\omega'(v, v') = \omega(v, v')$ if $(v, v') \in E$, and $\omega'(v, \mathtt{t}) = \omega'(\mathtt{t}, \mathtt{t}) = 0$. We show that $\mathsf{Val}_{\mathcal{G}}(v) < 0$ if and only if $\mathsf{Val}_{\mathcal{G}'}(v) = -\infty$.

In $\mathcal{G}'$, all values are different from $+\infty$, since Min plays at least every two steps, and has the capability to go to the target vertex with weight 0. Hence, letting $\mathcal{G}'' = \langle V', E', \omega', \mathbf{MP}\rangle$ the mean-payoff game on the weighted graph of $\mathcal{G}'$, by the previous direction, we have that for every vertex $v \in V'$, $\mathsf{Val}_{\mathcal{G}'}(v) = -\infty$ if and only if $\mathsf{Val}_{\mathcal{G}''}(v) < 0$.

To conclude, we prove that for all vertices $v \in V$, $\mathsf{Val}_{\mathcal{G}''}(v) < 0$ if and only if $\mathsf{Val}_{\mathcal{G}}(v) < 0$. If $\mathsf{Val}_{\mathcal{G}}(v) < 0$, by mapping the memoryless optimal strategies of $\mathcal{G}$ into $\mathcal{G}''$, we directly obtain that $\mathsf{Val}_{\mathcal{G}''}(v) \leqslant \mathsf{Val}_{\mathcal{G}}(v) < 0$. Reciprocally, if $\mathsf{Val}_{\mathcal{G}''}(v) < 0$, we can project a profile of memoryless optimal strategies over vertices of $\mathcal{G}$: the play obtained from v in $\mathcal{G}$ is then the projection of the play obtained from v in $\mathcal{G}''$, with the same cost. Hence, $\mathsf{Val}_{\mathcal{G}}(v) \leqslant \mathsf{Val}_{\mathcal{G}''}(v) < 0$. ◀

E Computing finite values in min-cost reachability games

We now prove the correction of Algorithm 1, as stated below.

▶ Proposition 5. If an MCR game $\mathcal{G} = \langle V, E, \omega, \mathbf{MCR}\rangle$ is given as input (possibly with values $+\infty$ or $-\infty$), Algorithm 1 outputs $\mathsf{Val}_{\mathcal{G}}$, after at most $(2|V| - 1)W|V| + 2|V|$ iterations.

Observe first that for all $v \in V$, for all $i \geqslant 0$, and for all strategies σ_{Max} and σ_{Min}:

$$\mathbf{MCR}^{\leqslant i}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) \geqslant \mathbf{MCR}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}}))\,.$$

Indeed, if the target vertex $\mathtt{t}$ is reached within i steps, then both plays will have the same payoff. Otherwise, $\mathbf{MCR}^{\leqslant i}(\mathsf{Play}(v, \sigma_{\mathsf{Min}}, \sigma_{\mathsf{Max}})) = +\infty$. Thus, for all $i \geqslant 1$ and $v \in V$:

$$\overline{\mathsf{Val}}^{\leqslant i}(v) \geqslant \overline{\mathsf{Val}}(v) = \mathsf{Val}(v)$$

which can be rewritten as

$$\overline{\mathsf{Val}}^{\leqslant i} \succcurlyeq \overline{\mathsf{Val}} = \mathsf{Val}\,.$$

Let us now consider the sequence $(\overline{\mathsf{Val}}^{\leqslant i})_{i \geqslant 0}$. We first give an alternative definition of this sequence permitting to show its convergence.

▶ **Lemma 6.** *For all $i \geqslant 1$, for all $v \in V$:*

$$\overline{\mathsf{Val}}^{\leqslant i}(v) = \begin{cases} \max\limits_{v' \in E(v)} \left(\omega(v, v') + \overline{\mathsf{Val}}^{\leqslant i-1}(v')\right) & \text{if } v \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\} \\ \min\limits_{v' \in E(v)} \left(\omega(v, v') + \overline{\mathsf{Val}}^{\leqslant i-1}(v')\right) & \text{if } v \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\} \\ 0 & \text{if } v = \mathtt{t} \end{cases}$$

Proof. The lemma can be established by showing that $\overline{\mathsf{Val}}^{\leqslant i}(v)$ is the value in a game played on a finite tree of depth i (i.e., by applying a backward induction). We adopt the following notation for labeled unordered trees. A leaf is denoted by (v), where $v \in V$ is the label of the leaf. A tree with root labeled by v and subtrees $A_1, \ldots, A_n$ is denoted by $(v, \{A_1, \ldots, A_n\})$. Then, for each $v \in V$ and $i \geqslant 0$, we define $A^i(v)$ as follows:

$$A^0(v) = (v)$$
$$\text{for all } i \geqslant 1 : A^i(v) = (v, \{A^{i-1}(v') \mid (v, v') \in E\})$$

Now, let us further label those trees by a value in $\mathbb{Z} \cup \{+\infty\}$ thanks to the function λ. For all tree of the form $A^0(v) = (v)$, we let

$$\lambda(A^0(v)) = \begin{cases} 0 & \text{if } v = \mathtt{t} \\ +\infty & \text{if } v \neg \mathtt{t} \end{cases}$$

For all tree of the form $A^i(v) = (v, \{A^{i-1}(v_1), \dots, A^{i-1}(v_m)\})$ (for some $i \geqslant 1$), we let

$$\lambda(A^i(v)) = \begin{cases} \max\limits_{1 \leqslant j \leqslant m} \big(\omega(v, v_j) + \lambda(A^{i-1}(v_j))\big) & \text{if } v \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\} \\ \min\limits_{1 \leqslant j \leqslant m} \big(\omega(v, v_j) + \lambda(A^{i-1}(v_j))\big) & \text{if } v \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\} \\ 0 & \text{if } v = \mathtt{t} \end{cases} \tag{1}$$

Clearly, for all $v \in V$, for all $i \geqslant 0$, the branches of $A^i(v)$ correspond to all the possible finite plays $\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})[i]$, i.e., there is a branch for each possible strategy profile $(\sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})$. Thus, $\lambda(A^i(v)) = \overline{\mathsf{Val}}^{\leqslant i}(v)$ for all $i \geqslant 0$, which permits us to conclude from (1). ◀

We have just shown that for all $i \geqslant 1$, $\overline{\mathsf{Val}}^{\leqslant i} = \mathcal{F}(\overline{\mathsf{Val}}^{\leqslant i-1})$, so that $x_i = \overline{\mathsf{Val}}^{\leqslant i}$ for all $i \geqslant 0$ (with x_i defined in page 8).

▶ Remark. Notice that, at this point, it would not be too difficult to show that Val is a fixed point of operator $\mathcal{F}$. However, it is more difficult to show that it is the greatest fixed point of $\mathcal{F}$, and to deduce directly properties over optimal strategies in the min-cost reachability game. Instead, we use the sequence $(\overline{\mathsf{Val}}^{\leqslant i})_{i \geqslant 0}$ to obtain more interesting results on Val.

We now study refined properties of the sequence $(\overline{\mathsf{Val}}^{\leqslant i})_{i \geqslant 0}$, namely its stationarity and the speed of its convergence. We start by characterizing how $\overline{\mathsf{Val}}^{\leqslant i}$ evolves over the first $|V| + 1$ steps. The next lemma states that, for each node v, the sequence $\overline{\mathsf{Val}}^{\leqslant 0}(v), \overline{\mathsf{Val}}^{\leqslant 1}(v), \dots, \overline{\mathsf{Val}}^{\leqslant i}(v), \dots, \overline{\mathsf{Val}}^{\leqslant |V|}(v)$ is of the form

$$\underbrace{+\infty, +\infty, \dots, +\infty}_{k \text{ times}}, a_k, a_{k+1}, \dots, a_{|V|}$$

where k is the step at which v has been added to the attractor, and each value a_i is finite and bounded:

▶ **Lemma 7.** *Let $v \in V$ be a vertex and let $0 \leqslant k \leqslant |V|$ be such that $v \in \mathsf{Attr}_k(\{\mathtt{t}\}) \setminus \mathsf{Attr}_{k-1}(\{\mathtt{t}\})$ (assuming $\mathsf{Attr}_{-1}(\{\mathtt{t}\}) = \emptyset$). Then, for all $0 \leqslant j \leqslant |V|$: (i) $j < k$ implies $\overline{\mathsf{Val}}^{\leqslant j}(v) = +\infty$ and (ii) $j \geqslant k$ implies $\overline{\mathsf{Val}}^{\leqslant j}(v) \leqslant jW$.*

Proof. We prove the property for all vertices v, by induction on j.

Base case: $j = 0$. We consider two cases. Either $v = \mathtt{t}$. In this case, $k = 0$, and we must show that $\overline{\mathsf{Val}}^{\leqslant 0}(v) \leqslant 0 \times W = 0$, which is true by definition of $\overline{\mathsf{Val}}^{\leqslant 0}$. Or $v \neq \mathtt{t}$. In this case, $k > 0$, and we must show that $\overline{\mathsf{Val}}^{\leqslant 0}(v) = +\infty$, which is true again by definition of $\overline{\mathsf{Val}}^{\leqslant 0}$.

Inductive case: $j = \ell \geqslant 1$. Let us assume that the lemma holds for all v, for all values of j up to $\ell - 1$, and let us show that it holds for all v, and for $j = \ell$. Let us fix a vertex v, and its associated value k. We consider two cases.

1. First, assume $k > \ell$. In this case, we must show that $\overline{\mathsf{Val}}^{\leqslant \ell}(v) = +\infty$. We consider again two cases:
   a. If $v \in V_{\mathsf{Min}}$, then none of its successors belong to $\mathsf{Attr}_{\ell-1}(\{\mathtt{t}\})$, otherwise, v would be in $\mathsf{Attr}_\ell(\{\mathtt{t}\})$, by definition of the attractor, and we would have $k \leqslant \ell$. Hence, by induction hypothesis, $\overline{\mathsf{Val}}^{\leqslant \ell-1}(v') = +\infty$ for all v' such that $(v, v') \in E$. Thus:
   $$\begin{aligned} \overline{\mathsf{Val}}^{\leqslant \ell}(v) &= \min_{(v,v') \in E} \Big(\omega(v, v') + \overline{\mathsf{Val}}^{\leqslant \ell-1}(v')\Big) && \text{(Lemma 6)} \\ &= +\infty \end{aligned}$$
   b. If $v \in V_{\mathsf{Max}}$, then at least one successor of v does not belong to $\mathsf{Attr}_{\ell-1}(\{\mathtt{t}\})$, otherwise, v would be in $\mathsf{Attr}_\ell(\{\mathtt{t}\})$, by definition of the attractor, and we would have $k \leqslant \ell$. Hence, by induction hypothesis, there exists v' such that $(v, v') \in E$ and $\overline{\mathsf{Val}}^{\leqslant \ell-1}(v') = +\infty$. Thus:
   $$\begin{aligned} \overline{\mathsf{Val}}^{\leqslant \ell}(v) &= \max_{(v,v') \in E} \Big(\omega(v, v') + \overline{\mathsf{Val}}^{\leqslant \ell-1}(v')\Big) && \text{(Lemma 6)} \\ &= +\infty \end{aligned}$$

2. Second, assume $k \leqslant \ell$. In this case, we must show that $\overline{\mathsf{Val}}^{\leqslant \ell}(v) \leqslant \ell W$. As in the previous item, we consider two cases:

a. In the case where $v \in V_{\mathsf{Min}}$, we let $\overline{v}$ be a vertex such that $\overline{v} \in \mathsf{Attr}_{k-1}(\{\mathtt{t}\})$ and $(v, \overline{v}) \in E$. Such a vertex exists by definition of the attractor. By induction hypothesis, $\overline{\mathsf{Val}}^{\leqslant \ell-1}(\overline{v}) \leqslant \ell W$. Then:

$$\begin{aligned}
\overline{\mathsf{Val}}^{\leqslant \ell}(v) &= \min_{(v,v')\in E} \left(\omega(v,v') + \overline{\mathsf{Val}}^{\leqslant \ell-1}(v')\right) && \text{(Lemma 6)}\\
&\leqslant \omega(v,\overline{v}) + \overline{\mathsf{Val}}^{\leqslant \ell-1}(\overline{v}) && ((v,\overline{v}) \in E)\\
&\leqslant \omega(v,\overline{v}) + (\ell-1)W && \text{(Ind. Hyp.)}\\
&\leqslant W + (\ell-1)W\\
&= \ell W
\end{aligned}$$

b. In the case where $v \in V_{\mathsf{Max}}$, we know that all successors v' of v belong to $\mathsf{Attr}_{k-1}(\{\mathtt{t}\})$ by definition of the attractor. By induction hypothesis, for all successors v' of v: $\overline{\mathsf{Val}}^{\leqslant \ell-1}(v') \leqslant \ell W$.Hence:

$$\begin{aligned}
\overline{\mathsf{Val}}^{\leqslant \ell}(v) &= \max_{(v,v')\in E} \left(\omega(v,v') + \overline{\mathsf{Val}}^{\leqslant \ell-1}(v')\right) && \text{(Lemma 6)}\\
&\leqslant \max_{(v,v')\in E} (W + (\ell-1)W) && \text{(Ind. Hyp.)}\\
&= \ell W\,.
\end{aligned}$$

◀

In particular, this allows us to conclude that, after $|V|$ steps, all values are bounded by $|V|W$:

▶ **Corollary 8.** *For all $v \in V$: $\overline{\mathsf{Val}}^{\leqslant |V|}(v) \leqslant |V|W$.*

The next step is to show that the sequence stabilises after a bounded number of steps:

▶ **Lemma 9.** *The sequence $\overline{\mathsf{Val}}^{\leqslant 0}, \ldots, \overline{\mathsf{Val}}^{\leqslant i}, \ldots$ stabilises after at most $(2|V|-1)W|V| + |V|$ steps.*

Proof. We first show that if Min can secure, from some vertex v, a payoff less than $-(|V|-1)W$, i.e., $\mathsf{Val}(v) < -(|V|-1)W$, then it can secure an arbitrarily small payoff from that vertex, i.e., $\mathsf{Val}(v) = -\infty$, which contradicts our hypothesis that the value is finite. Hence, let us suppose that there exists a strategy σ_{Min} for Min such that $\mathsf{Val}(v, \sigma_{\mathsf{Min}}) < -(|V|-1)W$. Let $\mathcal{G}'$ be the mean-payoff game studied in Proposition 4. We will show that $\mathsf{Val}_{\mathcal{G}'}(v) < 0$, which permits to conclude that $\mathsf{Val}_{\mathcal{G}}(v) = -\infty$. Let σ_{Max} be a memoryless strategy of Max. By hypothesis, we know that $\mathbf{MCR}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) < -(|V|-1)W$. This ensures the existence of a cycle with negative cost in the play $\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})$: otherwise, we could iteratively remove every possible nonnegative cycle of the finite play before reaching $\mathtt{t}$ (hence reducing the cost of the play) and obtain a play without cycles before reaching $\mathtt{t}$ with a cost less than $-(|V|-1)W$, which is impossible (since it should be of length at most $|V|-1$ to cross at most one occurrence of each vertex). Consider the first negative cycle in the play. After the first occurrence of the cycle, we let Min choose its actions like in the cycle. By this way, we can construct another strategy σ'_{Min} for Min, verifying that for every memoryless strategy σ_{Max} of Max, we have $\mathbf{MP}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma'_{\mathsf{Min}}))$ being the weight of the negative cycle in which the play finishes. Since for mean-payoff games, memoryless strategies are sufficient for Max, we deduce that $\mathsf{Val}_{\mathcal{G}'}(v) < 0$.

This reasoning permits to prove that at every step i, $\overline{\mathsf{Val}}^{\leqslant i}(v) \geqslant \mathsf{Val}(v) \geqslant -(|V|-1)W + 1$ for all vertices v. Recall from Corollary 8 that, after $|V|$ steps in the sequence, all vertices are assigned a value smaller that $|V|W$. Moreover, we know that the sequence is non-increasing. In summary, for all $k \geqslant 0$ and for all vertices v:

$$-(|V|-1)W + 1 \leqslant \overline{\mathsf{Val}}^{\leqslant |V|+k}(v) \leqslant |V|W$$

Hence, in the worst case a strictly decreasing sequence will need $(2|V|-1)W|V|$ steps to reach the lowest possible value where all vertices are assigned $-(|V|-1)W+1$ from the highest possible value where all vertices are assigned $|V|W$. Thus, taking into account the $|V|$ steps to reach a finite value on all vertices, the sequence stabilises in at most $(2|V|-1)W|V| + |V|$ steps. ◀

Let us thus denote by $\overline{\mathsf{Val}}^{\leqslant}$ the value obtained when the sequence $(\overline{\mathsf{Val}}^{\leqslant i})_{i\geqslant 0}$ stabilizes. We know from a previous discussion that $\overline{\mathsf{Val}}^{\leqslant}$ is the greatest fixed point of operator $\mathcal{F}$. We are now ready to prove that this value is the actual value of the game:

▶ **Lemma 10.** *For all min-cost reachability game:* $\overline{\mathsf{Val}}^{\leqslant} = \mathsf{Val}$.

Proof. We already know that $\overline{\mathsf{Val}}^{\leqslant} \succcurlyeq \mathsf{Val}$. Let us show that $\overline{\mathsf{Val}}^{\leqslant} \preccurlyeq \mathsf{Val}$. Let $v \in V$ be a vertex. Since $\mathsf{Val}(v)$ is finite integer, there exists a strategy σ_{Min} for Min that realises this value, i.e.,

$$\mathsf{Val}(v) = \sup_{\sigma_{\mathsf{Max}}} \mathbf{MCR}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) .$$

Notice that this holds because the values are integers, enducing that the infimum in the definition of $\overline{\mathsf{Val}}(v) = \mathsf{Val}(v)$ is indeed reached.

Let us build a tree $A_{\sigma_{\mathsf{Min}}}$ unfolding all possible plays from v against σ_{Min}. $A_{\sigma_{\mathsf{Min}}}$ has a root labeled by v. If a tree node is labeled by a vertex v of Min, this tree node has a unique child labeled by $\sigma_{\mathsf{Min}}(v)$. If a tree node is labeled by a vertex v of Max, this tree node has one child per successor v' of v in the graph, labeled by v'. We proceed this way until we encounter a node labeled by a vertex from $\mathtt{t}$ in which case this node is a leaf. $A_{\sigma_{\mathsf{Min}}}$ is necessarily finite. Otherwise, by König's Lemma, it has one infinite branch that never reaches $\mathtt{t}$. From that infinite branch, one can extract a strategy σ_{Max} for Max such that $\mathbf{MCR}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) = +\infty$, hence $\mathsf{Val}(v) = +\infty$, which contradicts the hypothesis. Assume the tree has depth m. Then, $A_{\sigma_{\mathsf{Min}}}$ is a subtree of the tree A obtained by unfolding all possible plays up to length m (as in the proof of Lemma 6). In this case, it is easy to check that the value labeling the root of $A_{\sigma_{\mathsf{Min}}}$ after applying backward induction is larger than or equal to the value labeling the root of A after applying backward induction. The latter is $\mathsf{Val}(v)$ while the former is $\overline{\mathsf{Val}}^{\leqslant m}(v)$, by Lemma 6, so that $\mathsf{Val}(v) \geqslant \overline{\mathsf{Val}}^{\leqslant m}(v)$. Since the sequence is non-increasing, we finally obtain $\mathsf{Val}(v) \geqslant \overline{\mathsf{Val}}^{\leqslant}(v)$. ◀

As a corollary of this lemma, we obtain:

▶ **Corollary 11.** Val *is the greatest fixed point of* $\mathcal{F}$.

This permits to obtain a value iteration algorithm, described in Algorithm 1, that computes optimal values. Notice that we do not suppose that every vertex has a finite value, which is justified in the proof of Proposition 5 below proving the correctness of the algorithm, as well as its complexity. A crucial argument is given in the following lemma, following from the fact that Val is the greatest fixed point of $\mathcal{F}$:

▶ **Lemma 12.** *If the Kleene sequence* $(\mathcal{F}^i(x_0))_{i \geqslant 0}$ *is initiated with a vector of values* x_0 *that is greater or equal to the optimal value vector* Val*, then the sequence converges at least as fast as before towards the optimal value vector.*

Proof of Proposition 5. Let us first suppose that values of every vertices are finite. Then, we can easily prove by induction that at the beginning of the jth step of the loop, X is equal to the vector $\overline{\mathsf{Val}}^{\leqslant j}$, and that the condition of line 7 has never been fulfilled. Hence, by Lemma 9, after at most $(2|V|-1)W|V| + |V|$ iterations, all values are found correctly in that case.

Suppose now that there exist vertices with value $+\infty$. Those vertices will remain at their initial value $+\infty$ during the whole computation, and hence do not interfere with the rest of the computation.

Finally, consider that the game contains vertices with value $-\infty$. We know that optimal values of vertices of values different from $-\infty$ are at least $-(|V|-1)W + 1$ so that, if the value of a vertex reaches an integer below $-(|V|-1)W$, we are sure that its value is indeed $-\infty$, which proves correct the line 7 of the algorithm. This update may cost at most one step per vertices, which in total adds at most $|V|$ iterations. Moreover, by Lemma 12, dropping the value to $-\infty$ does not harm the correction for the other vertices (it may only speed the convergence of their values). ◀

F Computing optimal strategies in min-cost reachability games

This section is devoted to the formal construction of optimal strategies in MCR games, as explained briefly at the end of Section 3.

F.1 Strategies of Min

We have already seen an example in Fig. 2 of a game where Min may need memory in an optimal strategy, i.e., where Min has an optimal strategy, but no memoryless optimal strategies. Reciprocally, as a consequence of the previous work, we first show that, for vertices with finite value, Min has always a finite-memory optimal strategy.

▶ Proposition 13. In all min-cost reachability game with only vertices of finite values, Min has a finite-memory optimal strategy.

Proof. We explain how to reconstruct from the fixpoint computation an optimal strategy σ^*_{Min} for Min. Let k be the integer such that $\overline{\mathsf{Val}}^{\leqslant k+1} = \overline{\mathsf{Val}}^{\leqslant k}$. For every step $\ell \leqslant k$ of the fixpoint computation, we define a strategy $\sigma^\ell_{\mathsf{Min}}$. For every finite play π ending in a vertex v of Min, if the length of π is $i < \ell$, we let

$$\sigma^\ell_{\mathsf{Min}}(\pi) = \operatorname*{argmin}_{v' \in E(v)} \left(\omega(v,v') + \overline{\mathsf{Val}}^{\leqslant \ell-i-1}(v')\right),$$

and otherwise, we let

$$\sigma^\ell_{\mathsf{Min}}(\pi) = \operatorname*{argmin}_{v' \in E(v)} \left(\omega(v,v') + \overline{\mathsf{Val}}^{\leqslant 0}(v')\right).$$

Notice that the argmin operator may select any possible vertex as long as it selects one with minimum value. We also let $\sigma^\ell_{\mathsf{Min}} = \sigma^k_{\mathsf{Min}}$ for every $\ell > k$, as well as $\sigma^*_{\mathsf{Min}} = \sigma^k_{\mathsf{Min}}$. Notice that σ^*_{Min} is a finite-memory strategy, since it only requires to know the last vertex and the length of the prefix up to k.

We now prove that $\mathsf{Val}(v, \sigma^*_{\mathsf{Min}}) = \sup_{\sigma_{\mathsf{Max}}} \mathbf{MCR}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^*_{\mathsf{Min}})) \leqslant \overline{\mathsf{Val}}^{\leqslant}(v)$ for all vertices v, which proves that σ^*_{Min} is an optimal strategy since $\overline{\mathsf{Val}}^{\leqslant}(v) = \mathsf{Val}(v)$ by Lemma 10. To do so, we first show by induction on ℓ that

$$\mathbf{MCR}^{\leqslant \ell}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^\ell_{\mathsf{Min}})) \leqslant \overline{\mathsf{Val}}^{\leqslant \ell}(v) \tag{2}$$

holds for every strategy σ_{Max} of Max and $v \in V$. This permits to conclude since, from $\sigma^*_{\mathsf{Min}} = \sigma^\ell_{\mathsf{Min}}$ for every $\ell \geqslant k$, we can deduce:

$$\begin{aligned}\lim_{\ell\to\infty} \sup_{\sigma_{\mathsf{Max}}} \mathbf{MCR}^{\leqslant \ell}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^\ell_{\mathsf{Min}})) &= \lim_{\ell\to\infty} \sup_{\sigma_{\mathsf{Max}}} \mathbf{MCR}^{\leqslant \ell}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^*_{\mathsf{Min}})) \\ &= \sup_{\sigma_{\mathsf{Max}}} \mathbf{MCR}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^*_{\mathsf{Min}}))\end{aligned}$$

and $\overline{\mathsf{Val}}^{\leqslant i}$ is a stationary sequence converging towards $\overline{\mathsf{Val}}^{\leqslant}$. The proof by induction goes as follows. In case $\ell = 0$, either $v = \mathtt{t}$ and both terms of (2) are equal to 0, or $v \neq \mathtt{t}$ and both terms of (2) are equal to $+\infty$. Supposing now that the property holds for an index ℓ, let us prove it for $\ell + 1$. For that, we consider a strategy σ_{Max} of Max. In case $v = \mathtt{t}$, we have

$$\mathbf{MCR}^{\leqslant \ell}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^\ell_{\mathsf{Min}})) = 0 = \overline{\mathsf{Val}}^{\leqslant \ell}(v).$$

We now consider the case $v \neq \mathtt{t}$. Let v' be the second vertex in $\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^{\ell+1}_{\mathsf{Min}})$. From the definition of $\sigma^{\ell+1}_{\mathsf{Min}}$, $\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^{\ell+1}_{\mathsf{Min}})[\ell+1]$ is the concatenation of v and $\mathsf{Play}(v', \sigma_{\mathsf{Max}}, \sigma^\ell_{\mathsf{Min}})[\ell]$. Hence,

$$\mathbf{MCR}^{\leqslant \ell+1}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^{\ell+1}_{\mathsf{Min}})) = \omega(v,v') + \mathbf{MCR}^{\leqslant \ell}(\mathsf{Play}(v', \sigma_{\mathsf{Max}}, \sigma^\ell_{\mathsf{Min}})).$$

By induction hypothesis, we obtain that

$$\mathbf{MCR}^{\leqslant \ell+1}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^{\ell+1}_{\mathsf{Min}})) \leqslant \omega(v,v') + \overline{\mathsf{Val}}^{\leqslant \ell}(v'). \tag{3}$$

Now, consider the two following cases.

- If $v \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\}$, we have $v' = \sigma^{\ell+1}_{\mathsf{Min}}(v)$, so that, in case $\ell + 1 \leqslant k$:

  $$\omega(v,v') + \overline{\mathsf{Val}}^{\leqslant \ell}(v') \leqslant \min_{v'' \in V \mid (v,v'') \in E} \left(\omega(v,v'') + \overline{\mathsf{Val}}^{\leqslant \ell}(v'')\right).$$

  Using (3) and Lemma 6, we obtain

  $$\mathbf{MCR}^{\leqslant \ell+1}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^{\ell+1}_{\mathsf{Min}})) \leqslant \overline{\mathsf{Val}}^{\leqslant \ell+1}(v).$$

  In case $\ell + 1 > k$, we indeed have $\overline{\mathsf{Val}}^{\leqslant \ell} = \overline{\mathsf{Val}}^{\leqslant \ell+1} = \overline{\mathsf{Val}}^{\leqslant k}$, so that we conclude similarly.
- If $v \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\}$, we have $v' = \sigma_{\mathsf{Max}}(v)$ and

  $$\omega(v,v') + \overline{\mathsf{Val}}^{\leqslant \ell}(v') \leqslant \max_{v'' \in V \mid (v,v'') \in E} \left(\omega(v,v'') + \overline{\mathsf{Val}}^{\leqslant \ell}(v'')\right).$$

  Once again using (3) and Lemma 6, we obtain

  $$\mathbf{MCR}^{\leqslant \ell+1}(\mathsf{Play}(v, \sigma_{\mathsf{Max}}, \sigma^{\ell+1}_{\mathsf{Min}})) \leqslant \overline{\mathsf{Val}}^{\leqslant \ell+1}(v).$$

This concludes the induction proof. ◀

Notice that the proof of Proposition 13 together with the statement of Lemma 9 imply that a memory of size pseudo-polynomial for the strategy of Min is sufficient. Before stating the result for Max, we informally refine this result in order to find a strategy of Min having more *structural properties*. It will be composed of two memoryless strategies σ^1_{Min} and σ^2_{Min}: the game will start with Min following σ^1_{Min}, and at some point, determined by the weight of the current finite play, Min will switch to strategy σ^2_{Min} which is an attractor strategy, i.e., a strategy that reaches the target in less than $|V|$ steps, regardless of the weights along this path. Intuitively the strategy σ^1_{Min} ensures either to reach the target with optimal value, or to go in cycles of negative weights. The only chance for Min of having a greater value than the optimal is to go infinitely through these cycles without reaching the target. But if it does so, the total-payoff will decrease and at some point the value will be so low, that the cost of calling the attractor strategy will leave the total-payoff smaller than the optimal value. Let us formalise this construction.

For the sake of exposure, we present the construction when all values are finite, but such construction can be applied with few changes when some vertices have value $-\infty$ or $+\infty$. We start by defining a memoryless strategy σ^1_{Min} that has some good properties (stated in Proposition 15). Let X^i denote the value of variable X after i iteration of the loop of Algorithm 1, and let $X^0(v) = +\infty$ for all $v \in V$. We have seen that the sequence $X^0 \succcurlyeq X^1 \succcurlyeq X^2 \succcurlyeq \cdots$ is stationary at some point, equal to Val. For all vertices $v \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\}$, let $i_v > 0$ be the first index such that $X^i(v) = \mathsf{Val}(v)$. Fix a vertex $v' \neq \mathtt{t}$ such that $X^i(v) = \omega(v, v') + X^{i-1}(v')$ (that exists by definition) and define $\sigma^1_{\mathsf{Min}}(v) = v'$. The following lemma states that the vertex v' already reached its final value at step $i - 1$.

▶ **Lemma 14.** *For all vertices $v \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\}$, $X^{i-1}(\sigma^1_{\mathsf{Min}}(v)) = \mathsf{Val}(\sigma^1_{\mathsf{Min}}(v))$.*

Proof. Let $v' = \sigma^1_{\mathsf{Min}}(v)$. By contradiction assume that $X^{i-1}(v') > \mathsf{Val}(v')$, Note that there exists $j > i$ such that $X^{j-1}(v') = \mathsf{Val}(v')$. By definition,

$$\begin{aligned}\mathsf{Val}(v) \leqslant X^j(v) \leqslant \omega(v, v') + X^{j-1}(v') = \omega(v, v') + \mathsf{Val}(v') \\ < \omega(v, v') + X^{i-1}(v') = X^i(v) = \mathsf{Val}(v),\end{aligned}$$

which raises a contradiction. ◀

We can state the properties of σ^1_{Min}: intuitively one can see it as an *almost perfect strategy*, in the sense that it is memoryless, if it reaches the target, then the value obtained is optimal, and if it does not reach the target then the total-payoff of the finite play will decrease as the game goes on. The only problem is one cannot ensure that we reach the target.

▶ Proposition 15. For all vertices v, and for all plays $\pi = v_1 v_2 \cdots$ starting in v and conforming to σ^1_{Min},

(1) if there exists $i < j$ such that $v_i = v_j$, then $\mathbf{TP}(v_i \cdots v_j) < 0$,
(2) if π reaches $\mathtt{t}$ then $\mathbf{TP}(\pi) \leqslant \mathsf{Val}(v)$.

Proof. Let us prove (1), take a cycle $v_i \cdots v_j$ with $v_j = v_i$. Notice that at least one vertex of this cycle belongs to Min, since, otherwise, Max would have a strategy to obtain a value $+\infty$ for vertex v_1, which contradicts the hypothesis on the game. Hence, for the sake of the explanation, we suppose that $v_i \in V_{\mathsf{Min}}$. Let us also suppose that i_{v_0} is maximal among $\{i_{v_\ell} \mid i \leqslant \ell < j, v_{i'} \in V_{\mathsf{Max}}\}$. The following extends straightforwardly to the case where this maximal vertex of Min is not v_i. We prove by induction over $i < \ell \leqslant j$ that

$$X^{i_{v_i}}(v_i) \geqslant \mathbf{TP}(v_i \cdots v_\ell) + X^{i_{v_i}-1}(v_\ell)\,.$$

The base case comes from the fact that since $v_i \in V_{\mathsf{Min}}$ we have $\sigma^1_{\mathsf{Min}}(v_i) = v_{i+1}$ thus $X^{i_{v_0}} = \omega(v_i, v_{i+1}) + X^{i_{v_0}-1}(v_{i+1})$. For the inductive case, let us consider $i < \ell < j$ such that $X^{i_{v_i}}(v_i) \geqslant \mathbf{TP}(v_i \cdots v_\ell) + X^{i_{v_i}-1}(v_\ell)$ and let us prove it for $\ell + 1$.

If $v_\ell \in V_{\mathsf{Min}}$, by definition of $X^{i_{v_i}}$, we have

$$X^{i_{v_0}}(v_\ell) = \max_{(v_\ell, v') \in E} \omega(v_\ell, v') + X^{i_{v_0}-1}(v') \geqslant \omega(v_\ell, v_{\ell+1}) + X^{i_{v_0}-1}(v_{\ell+1})\,.$$

In case $v_j \in V_{\mathsf{Min}}$, by maximality of i_{v_0}, we have

$$\begin{aligned}X^{i_{v_0}}(v_\ell) &= X^{i_{v_\ell}}(v_\ell) = \omega(v_\ell, v_{\ell+1}) + X^{i_{v_\ell}-1}(v_{\ell+1}) \\ &\geqslant \omega(v_\ell, v_{\ell+1}) + X^{i_{v_0}-1}(v_{\ell+1})\end{aligned}$$

using that the sequence $X^0, X^1, X^2, \dots$ is non-increasing.

Hence, in all cases, we have

$$X^{i_{v_0}}(v_\ell) \geqslant \omega(v_\ell, v_{\ell+1}) + X^{i_{v_0}-1}(v_{\ell+1})$$

Using again that $X^0, X^1, X^2, \dots$ is non-decreasing, we obtain

$$X^{i_{v_0}-1}(v_\ell) \geqslant X^{i_{v_0}}(v_\ell) \geqslant \omega(v_\ell, v_{\ell+1}) + X^{i_{v_0}-1}(v_{\ell+1})\,.$$

Injecting this into the induction hypothesis, we have

$$\begin{aligned} X^{i_{v_0}}(v_i) &\geqslant \mathbf{TP}(v_i \cdots v_\ell) + \omega(v_\ell, v_{\ell+1}) + X^{i_{v_0}-1}(v_{\ell+1}) \\ &= \mathbf{TP}(v_i \cdots v_{\ell+1}) + X^{i_{v_0}-1}(v_{\ell+1}) \end{aligned}$$

which concludes the proof by induction. In particular, for $\ell = j$, as $v_i = v_j$ we obtain that

$$X^{i_{v_0}}(v_i) \geqslant X^{i_{v_0}-1}(v_i) + \mathbf{TP}(v_i \cdots v_j)\,.$$

and as, by definition of i_{v_0}, we have $X^{i_{v_0}}(v_i) < X^{i_{v_0}-1}(v_i)$, we necessarily have $\mathbf{TP}(v_1 \cdots v_j) < 0$.

To prove (2) we decompose π as $\pi = v_1 \cdots v_k \mathtt{t}^\omega$ with for all i, $v_i \neq \mathtt{t}$. We prove by decreasing induction on i that $\mathbf{TP}(v_i \cdots v_k \mathtt{t}^\omega) \leqslant \mathsf{Val}(v)$. If $i = k+1$, $\mathbf{TP}(\mathtt{t}^\omega) = 0 = \mathsf{Val}(\mathtt{t})$. If $i \leqslant k$, by induction we have $\mathbf{TP}(v_{i+1} \cdots v_k \mathtt{t}^\omega) = \mathsf{Val}(v_{i+1})$ thus $\mathbf{TP}(v_i \cdots v_k \mathtt{t}^\omega) = \omega(v_i, v_{i+1}) + \mathsf{Val}(v_{i+1})$. If $v_i \in V_{\mathsf{Min}}$ then $v_{i+1} = \sigma^\star_{\mathsf{Min}}(v)$ and $\mathsf{Val}(v_i) = \omega(v_i, v_{i+1}) + \mathsf{Val}(v_{i+1}) = \mathbf{TP}(v_i \cdots v_k \mathtt{t}^\omega)$. If $v_i \in V_{\mathsf{Max}}$, then $\mathbf{TP}(v_i \cdots v_k \mathtt{t}^\omega) = \omega(v_i, v_{i+1}) + \mathsf{Val}(v_{i+1}) \leqslant \max_{(v_i, v') \in E} \omega(v_i, v_{i+1}) + \mathsf{Val}(v_{i+1}) = \mathsf{Val}(v_i)$. ◀

Next, let σ^2_{Min} be the memoryless strategy induced by the computation of the attractor: notice that it is possible to construct it directly from the value iteration computation by mapping a vertex v to one vertex from which v is first discovered (i.e., its value is first set to a real value different from $+\infty$). This strategy ensures to reach the target after at most $|V|$ steps, thus for all v, $\mathsf{Val}(v, \sigma^2_{\mathsf{Min}}) \leqslant W(|V|-1)$.

Before defining the strategy $\widetilde{\sigma_{\mathsf{Min}}}$, we introduce the notion of *switchable finite play* as follows. A finite play $v_1 \cdots v_k$ is switchable if $v_k \in V_{\mathsf{Max}}$ and $\mathbf{TP}(v_1 \cdots v_k) \leqslant \mathsf{Val}(v_1) - \mathsf{Val}(v_k, \sigma^2_{\mathsf{Min}})$. Intuitively a strategy is switchable, if by switching to the attractor strategy, we ensures to get an optimal value.

We define the strategy σ'_{Min} as follows: for all finite play $v_1 \cdots v_k$ with $v_k \in V_{\mathsf{Max}}$:

$$\sigma'_{\mathsf{Min}}(v_1 \cdots v_k) = \begin{cases} \sigma^2_{\mathsf{Min}}(v_k) & \text{if } v_1 \cdots v_i \text{ is switchable for some } i \\ \sigma^1_{\mathsf{Min}}(v_k) & \text{otherwise} \end{cases}$$

One can easily show that for all play $v_1 v_2 \cdots$ that conforms to σ'_{Min} either $v_1 v_2 \cdots$ conforms to σ^1_{Min} if no finite play is switchable, or there exists k such that $v_1 \cdots v_k$ conforms to σ^1_{Min}, $v_1 \cdots v_k$ is a switchable finite play, and $v_k v_{k+1} \cdots$ conforms to σ^2_{Min}. The following proposition states that the strategy σ'_{Min} is optimal.

▶ Proposition 16. For all v, $\mathsf{Val}(v, \sigma'_{\mathsf{Min}}) = \mathsf{Val}(v)$.

Proof. Let $v \in V$ and $\pi = v_1 v_2 \cdots \in \mathsf{Play}(v, \sigma'_{\mathsf{Min}})$, let us show that $\mathbf{MCR}(\pi) \leqslant \mathsf{Val}(v)$. Assume first that there exists k such that $v_1 \cdots v_k$ conforms to σ^1_{Min}, $v_1 \cdots v_k$ is a switchable finite play, and $v_k v_{k+1} \cdots$ conforms to σ^2_{Min}. Thus π reaches the target as $v_k v_{k+1} \cdots$ does, and $\mathbf{MCR}(\pi) = \mathbf{TP}(v_1 \cdots v_k) + \mathbf{MCR}(v_k v_{k+1} \cdots)$. As $v_1 \cdots v_k$ is a switchable finite play, we have $\mathbf{TP}(v_1 \cdots v_k) \leqslant \mathsf{Val}(v_1) - \mathsf{Val}(v_k, \sigma^2_{\mathsf{Min}})$, thus $\mathbf{MCR}(\pi) \leqslant \mathsf{Val}(v_1) - \mathsf{Val}(v_k, \sigma^2_{\mathsf{Min}}) + \mathbf{MCR}(v_k v_{k+1} \cdots) \leqslant \mathsf{Val}(v_1)$.

Assume now that $v_1 v_2 \cdots$ does not contains a switchable prefix and thus π conforms to σ^1_{Min}. If π reaches $\mathtt{t}$ then by Proposition 15, $\mathbf{MCR}(\pi) \leqslant \mathsf{Val}(v)$. To conclude, let us prove that π reaches $\mathtt{t}$, and by contradiction assume that this is not the case.

First we prove by induction on k, that ($\star$) for all play $v'_1 \cdots v'_k$ that conforms to σ^1_{Min} and does not reach $\mathtt{t}$, $\mathbf{TP}(v_1 \cdots v_k) \leqslant (|V|-1)W - \left\lfloor \frac{k}{|V|} \right\rfloor$.

If $k < |V|$ the results is straightforward. For the inductive case, let $k \geqslant |V|$ and $v'_1 \cdots v'_k$ that conforms to σ^1_{Min} and does not reach $\mathtt{t}$. Then as $k \geqslant |V|$, there exists $i < j$ such that $v_i = v_j$. Thus as $v'_1 \cdots v'_i v'_{j+1} \cdots v'_k$ is a play of size $k - |V|$ that conforms to σ^1_{Min} and does not reach $\mathtt{t}$, we know by induction hypothesis that $\mathbf{TP}(v'_1 \cdots v'_i v'_{j+1} \cdots v'_k) \leqslant (|V|-1)W - \left\lfloor \frac{k-|V|}{|V|} \right\rfloor = (|V|-1)W - \left\lfloor \frac{k}{|V|} \right\rfloor + 1$. Furthermore by Proposition 15, we have that $\mathbf{TP}(v'_i \cdots v'_j) \leqslant -1$. Thus

$$\begin{aligned} \mathbf{TP}(v'_1 \cdots v'_k) &= \mathbf{TP}(v'_1 \cdots v'_i) + \mathbf{TP}(v'_i \cdots v'_j) + \mathbf{TP}(v'_j \cdots v'_k) \\ &= \mathbf{TP}(v'_1 \cdots v'_i v'_{j+1} \cdots v'_k) + \mathbf{TP}(v'_1 \cdots v'_j) \leqslant (|V|-1)W - \left\lfloor \frac{k}{|V|} \right\rfloor, \end{aligned}$$

which concludes the proof of $(\star)$.

Now we can get back to raising a contradiction, and for that we show that there exists a switchable finite play in $v_1 v_2 \cdots$. Take k be the least index greater that $3|V|(|V|-1)W$ such that $v_k \in V_{\mathsf{Max}}$ (we know that there exists one, otherwise the vertices v_j with $j \geqslant 3|V|(|V|-1)W$ would not be in the attractor of $\mathtt{t}$). We have

$$\mathbf{TP}(v_1 \cdots v_k) \leqslant (|V|-1)W - \left\lfloor \frac{k}{|V|} \right\rfloor \leqslant -2(|V|-1)W \leqslant \mathsf{Val}(v_1) + \mathsf{Val}(v_k, \sigma^2_{\mathsf{Min}}).$$

Thus $v_1 \cdots v_k$ is a switchable prefix, which raises a contradiction. ◀

Notice that σ'_{Min} may be more easily implementable than a more general finite-memory strategy, in particular, we may encode the current total-payoff in binary, hence saving some space. We give in Algorithm 3 a way to compute strategies σ^1_{Min} and σ^2_{Min}.

F.2 Strategies of Max

While we have already shown that optimal strategies for Min might require memory, let us show that Max always has a *memoryless optimal strategy*. This asymmetry stems directly from the asymmetric definition of the game – while Min has the double objective of reaching $\mathtt{t}$ and minimising its cost, Max aims at avoiding $\mathtt{t}$, and if not possible, maximising the cost.

▶ Proposition 17. In all min-cost reachability game, Max has a memoryless optimal strategy.

Proof. For vertices with value $+\infty$, we already know a memoryless optimal strategy for Max, namely any strategy that remains outside the attractor of the target vertices. For vertices with value $-\infty$, all strategies are equally bad for Max.

We now explain how to define a memoryless optimal strategy σ^*_{Max} for Max in case of a graph containing only finite values. For every finite play π ending in a vertex $v \in V_{\mathsf{Max}}$ of Max, we let

$$\sigma^*_{\mathsf{Max}}(\pi) = \operatorname*{argmax}_{v' \in E(v)} \left(\omega(v, v') + \mathsf{Val}(v')\right).$$

This is clearly a memoryless strategy. Let us prove that it is optimal for Max, that is, for every vertex $v \in V$, and every strategy σ_{Min} of Min

$$\mathbf{MCR}(\mathsf{Play}(v, \sigma^*_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) \geqslant \mathsf{Val}(v).$$

In case $\mathsf{Play}(v, \sigma^*_{\mathsf{Max}}, \sigma_{\mathsf{Min}})$ does not reach the target set of vertices, the inequality holds trivially. Otherwise, we let $\mathsf{Play}(v, \sigma^*_{\mathsf{Max}}, \sigma_{\mathsf{Min}}) = v_0 v_1 \cdots v_\ell \cdots$ with ℓ the least position such that $v_\ell = \mathtt{t}$. If $\ell = 0$, i.e., $v = v_0 = \mathtt{t}$, we have

$$\mathbf{MCR}(\mathsf{Play}(v, \sigma^*_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) = 0 = \mathsf{Val}(v).$$

Otherwise, let us prove by induction on $0 \leqslant i \leqslant \ell$ that

$$\mathbf{MCR}(v_{\ell-i} \cdots v_\ell) \geqslant \mathsf{Val}(v_{\ell-i}).$$

This will permit to conclude since

$$\mathbf{MCR}(\mathsf{Play}(v, \sigma^*_{\mathsf{Max}}, \sigma_{\mathsf{Min}})) = \mathbf{MCR}(v_0 v_1 \cdots v_\ell) \geqslant \mathsf{Val}(v_0) = \mathsf{Val}(v).$$

The base case $i = 0$ corresponds to the previous case where the starting vertex is $\mathtt{t}$. Supposing that the property holds for index i, let us prove it for $i+1$. We have

$$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) = \omega(v_{\ell-i-1}, v_{\ell-i}) + \mathbf{MCR}(v_{\ell-i} \cdots v_\ell).$$

By induction hypothesis, we have

$$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) \geqslant \omega(v_{\ell-i-1}, v_{\ell-i}) + \mathsf{Val}(v_{\ell-i}). \tag{4}$$

We now consider two cases:

- If $v_{\ell-i-1} \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\}$, then $v_{\ell-i} = \sigma^*_{\mathsf{Max}}(v_0 v_1 \cdots v_{\ell-i-1})$, so that by definition of σ^*_{Max}:

  $$\omega(v_{\ell-i-1}, v_{\ell-i}) + \mathsf{Val}(v_{\ell-i}) = \max_{v' \in V \mid (v_{\ell-i-1}, v') \in E} \left(\omega(v_{\ell-i-1}, v') + \mathsf{Val}(v')\right).$$

  Using Corollary 11 and (4), we obtain

  $$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) \geqslant \mathsf{Val}(v_{\ell-i-1}).$$

- If $v_{\ell-i-1} \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\}$, then

$$\omega(v_{\ell-i-1}, v_{\ell-i}) + \mathsf{Val}(v_{\ell-i}) \geqslant \min_{v' \in V \mid (v_{\ell-i-1}, v') \in E} \big(\omega(v_{\ell-i-1}, v') + \mathsf{Val}(v')\big)\,.$$

Once again using Corollary 11 and (4), we obtain

$$\mathbf{MCR}(v_{\ell-i-1} \cdots v_\ell) \geqslant \mathsf{Val}(v_{\ell-i-1})\,.$$

This concludes the proof. ◀

This strategy $\sigma^\star_{\mathsf{Max}}$ can directly be computed along the execution of the value iteration algorithm. This is done in Algorithm 3.

Algorithm 3: Computation of optimal strategy for both players in value iteration algorithm for min-cost reachability games

Input: min-cost reachability game $\langle V, E, \omega, \mathbf{MCR}\rangle$, W greatest weight in absolute value in the arena

$\mathsf{X}(\mathtt{t}) := 0$
foreach $v \in V \setminus \{\mathtt{t}\}$ **do** $\mathsf{X}(v) := +\infty$
repeat
  $\mathsf{X}_{pre} := \mathsf{X}$
  foreach $v \in V_{\mathsf{Max}} \setminus \{\mathtt{t}\}$ **do**
    $\mathsf{X}(v) := \max_{v' \in E(v)} \big(\omega(v, v') + \mathsf{X}_{pre}(v')\big)$
    if $\mathsf{X}(v) \neq \mathsf{X}_{pre}(v)$ **then** $\sigma^*_{\mathsf{Max}}(v) := \operatorname{argmax}_{v' \in E(v)} \big(\omega(v, v') + \mathsf{X}_{pre}(v')\big)$
  foreach $v \in V_{\mathsf{Min}} \setminus \{\mathtt{t}\}$ **do**
    $\mathsf{X}(v) := \min_{v' \in E(v)} \big(\omega(v, v') + \mathsf{X}_{pre}(v')\big)$
    if $\mathsf{X}(v) \neq \mathsf{X}_{pre}(v)$ **then**
      $\sigma^1_{\mathsf{Min}}(v) := \operatorname{argmax}_{v' \in E(v)} \big(\omega(v, v') + \mathsf{X}_{pre}(v')\big)$
      if $\mathsf{X}_{pre}(v) = +\infty$ **then** $\sigma^2_{\mathsf{Min}}(v) = \sigma^1_{\mathsf{Min}}(v)$
  foreach $v \in V \setminus \{\mathtt{t}\}$ **do**
    if $\mathsf{X}(v) < -(|V|-1)W$ **then** $\mathsf{X}(v) := -\infty$
until $\mathsf{X} = \mathsf{X}_{pre}$
return X

G Reduction of total-payoff games to min-cost reachability games

This section is devoted to the proof of Proposition 2.

For that purpose, we must relate paths in games $\mathcal{G}$ and $\mathcal{G}^n$: with each finite path in $\mathcal{G}^n$, we associate a finite path in $\mathcal{G}$, obtained by looking at the sequence of vertices of V appearing inside the vertices of the finite play. Formally, the *projection* of a finite path π is the sequence $\mathsf{proj}(\pi)$ of vertices of $\mathcal{G}$ inductively defined by $\mathsf{proj}(\varepsilon) = \varepsilon$ and for all finite path π, $v \in V$ and $1 \leqslant j \leqslant n$:

$$\mathsf{proj}((\mathtt{in}, v, j)\pi) = \mathsf{proj}(\pi)\,, \qquad \mathsf{proj}((\mathtt{ex}, v, j)\mathtt{t}\pi) = v\,, \qquad \mathsf{proj}((\mathtt{ex}, v, j)) = \varepsilon\,,$$
$$\mathsf{proj}((\mathtt{ex}, v, j+1)(v, j)\pi) = \mathsf{proj}((v, j)\pi) = v\,\mathsf{proj}(\pi)\,.$$

In particular, notice that in the case of a play with prefix $(\mathtt{ex}, v, j)\mathtt{t}$, the rest of the play is entirely composed of target vertices $\mathtt{t}$, since $\mathtt{t}$ is a sink state. For instance, the projection of the finite play

$$(v_1, 3)(\mathtt{in}, v_2, 3)(\mathtt{ex}, v_2, 3)(v_2, 2)(\mathtt{in}, v_3, 2)(v_3, 2)(\mathtt{in}, v_3, 2)(\mathtt{ex}, v_3, 2)\mathtt{t}$$

of the game $\mathcal{G}^3$ of Fig. 3 is given by $v_1v_2v_3v_3$.

The following lemma relates plays of $\mathcal{G}^n$ with their projection in $\mathcal{G}$, comparing their total-payoff.

▶ **Lemma 18.** *The projection mapping satisfies the following properties.*

1. *If π is a finite play in $\mathcal{G}^n$ then $\mathsf{proj}(\pi)$ is a finite play in $\mathcal{G}$.*
2. *If π is a play in $\mathcal{G}^n$ that does not reach the target, then $\mathsf{proj}(\pi)$ is a play in $\mathcal{G}$.*
3. *For all finite play π, $\mathbf{TP}(\pi) = \mathbf{TP}(\mathsf{proj}(\pi))$.*

Proof. The proof of 2 is a direct consequence of 1. To each vertex $w \in V^n \setminus \{\mathtt{t}\}$, we associate a vertex $f(w)$ as follows:

$$f(v,j) = f(\mathtt{in}, v, j) = f(\mathtt{ex}, v, j) = v\,.$$

Then notice that if $(w, w') \in E^n$ with $w, w' \neq \mathtt{t}$, then either $f(w) = f(w')$ or $(f(w), f(w')) \in E$. We now prove 1 and 3 inductively on the size of the finite play $\pi = w_1 \cdots w_k$ of $\mathcal{G}^n$, along with the fact that

$$\text{4. if } \mathsf{proj}(\pi) \neq \emptyset \text{ and } w_1 \neq \mathtt{t} \text{ then the first vertex of } \mathsf{proj}(\pi) \text{ is } f(w_1).$$

If $k = 0$, then $\pi = \mathsf{proj}(\pi) = \varepsilon$ are finite plays with the same total-payoff. If $k = 1$, either $\mathsf{proj}(\pi) = \varepsilon$ or $\pi = (v, j)$ and $\mathsf{proj}(\pi) = v$: in both cases, the properties hold trivially. Otherwise, $k \geqslant 2$ and we distinguish several possible prefixes:

- If $\pi = (\mathtt{in}, v, j)\pi'$, then $\mathsf{proj}(\pi) = \mathsf{proj}(\pi')$. Hence, 1 holds by induction hypothesis. If $\mathsf{proj}(\pi)$ is non-empty, so is $\mathsf{proj}(\pi')$. Moreover, the first vertex of π' is either (v, j) or $(\mathtt{ex}, v, j)$, so that we can show 4 by induction hypothesis. Finally, the previous remark shows that the first edge of π has necessarily weight 0, so that, $\mathbf{TP}(\pi) = \mathbf{TP}(\pi')$, and 3 also holds by induction hypothesis.
- If $\pi = (v, j)\pi'$, then $\mathsf{proj}(\pi) = v\,\mathsf{proj}(\pi')$ so that 4 holds directly. Moreover, π' is a non-empty finite play so that $\pi' = (\mathtt{in}, v', j)\pi''$ with $(v, v') \in E$, and $\mathsf{proj}(\pi') = \mathsf{proj}(\pi'')$. By induction, $\mathsf{proj}(\pi')$ is a finite play in $\mathcal{A}$, and it starts with v' (by 4). Since $(v, v') \in E$, this shows that $\mathsf{proj}(\pi)$ is a finite play. Moreover, $\mathbf{TP}(\pi) = \omega^n((v, j), (\mathtt{in}, v', j)) + \mathbf{TP}(\pi') = \omega(v, v') + \mathbf{TP}(\pi')$. By induction hypothesis, we have $\mathbf{TP}(\pi') = \mathbf{MCR}(\mathsf{proj}(\pi'))$. Moreover, $\mathbf{MCR}(\mathsf{proj}(\pi)) = \omega(v, v') + \mathbf{MCR}(\mathsf{proj}(\pi'))$ which concludes the proof of 3.
- If $\pi = (\mathtt{ex}, v, j)(v, j-1)\pi'$ then $\mathsf{proj}(\pi) = v\,\mathsf{proj}(\pi') = \mathsf{proj}((v, j-1)\pi')$: this allows us to conclude directly by using the previous case.
- Otherwise, $\pi = (\mathtt{ex}, v, j)\mathtt{t}\pi'$, and then $\mathsf{proj}(\pi) = v$ is a finite play with total-payoff 0, like π, and 4 holds trivially.

◀

The next lemma states that when playing memoryless strategies, one can bound the total-payoff of all finite plays.

▶ **Lemma 19.** *Let $v \in V$, and σ_{Max} (respectively, σ_{Min}) be a memoryless strategy for* Max *(respectively,* Min*) in the total-payoff game $\mathcal{G}$, such that $\mathsf{Val}(v, \sigma_{\mathsf{Max}}) \neq -\infty$ (respectively, $\mathsf{Val}(v, \sigma_{\mathsf{Min}}) \neq +\infty$). Then for all finite play π conforming to σ_{Max} (respectively, to σ_{Min}), $\mathbf{TP}(\pi) \geqslant -(|V|-1)W$ (respectively, $\mathbf{TP}(\pi) \leqslant (|V|-1)W$).*

Proof. We prove the part for Min, the other case is similar. The proof proceeds by induction on the size of a partial play $\pi = v_1 \cdots v_k$ with $v_1 = v$. If $k \leqslant |V|$ then $\mathbf{TP}(\pi) = \sum_{i=1}^{k-1} \omega(v_i, v_{i+1}) \leqslant (k-1)W \leqslant (|V|-1)W$. If $k \geqslant |V|+1$ then there exists $i < j$ such that $v_i = v_j$. Assume by contradiction that $\mathbf{TP}(v_i \cdots v_j) > 0$. Then the play $\pi' = v_1 \cdots v_i \cdots v_j (v_{i+1} \cdots v_j)^\omega$ conforms to σ_{Min} and $\mathbf{TP}(\pi') = +\infty$ which contradicts $\mathsf{Val}(v, \sigma_{\mathsf{Min}}) \neq +\infty$. Therefore $\mathbf{TP}(v_i \cdots v_j) \leqslant 0$. We have $\mathbf{TP}(\pi) = \mathbf{TP}(v_1 \cdots v_i) + \mathbf{TP}(v_i \cdots v_j) + \mathbf{TP}(v_{j+1} \cdots v_k)$, and since $v_i = v_j$, $v_1 \cdots v_i v_{j+1} \cdots v_k$ is a finite play starting from v that conforms to σ_{Min}, and by induction hypothesis $\mathbf{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) \leqslant (|V|-1)W$. Then $\mathbf{TP}(\pi) = \mathbf{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) + \mathbf{TP}(v_i \cdots v_j) \leqslant \mathbf{TP}(v_1 \cdots v_i v_{j+1} \cdots v_k) \leqslant (|V|-1)W$. ◀

This permits to bound the finite values $\mathsf{Val}(v)$ of vertices v of the game:

▶ **Corollary 20.** *For all $v \in V$, $\mathsf{Val}(v) \in \big[-(|V|-1)W, (|V|-1)W\big] \uplus \{-\infty, +\infty\}$.*

Proof. From [10], we know that total-payoff games are positionally determined, i.e., there exists two memoryless strategies $\sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}}$ such that for all v, $\mathsf{Val}(v) = \mathsf{Val}(v, \sigma_{\mathsf{Max}}) = \mathsf{Val}(v, \sigma_{\mathsf{Min}})$. Assume that $\mathsf{Val}(v) \notin \{-\infty, +\infty\}$. Then since $\mathsf{Val}(v, \sigma_{\mathsf{Min}}) = \mathsf{Val}(v) \neq -\infty$, Lemma 19 shows that all finite play π that conforms to σ_{Min} verifies $\mathbf{TP}(\pi) \geqslant -(|V|-1)W$, therefore $\mathsf{Val}(v) \geqslant -(|V|-1)W$. One can similarly prove that $\mathbf{TP}(v) \leqslant (|V|-1)W$. ◀

We now compare values in both games. A first lemma shows, in particular, that $\mathsf{Val}_{\mathcal{G}^n}(v, n) \leqslant \mathsf{Val}_{\mathcal{G}}(v)$, in case $\mathsf{Val}_{\mathcal{G}}(v) \neq +\infty$.

▶ **Lemma 21.** *For all $m \in \mathbb{Z}$, $v \in V$, and $n \geqslant 1$, if $\mathsf{Val}_{\mathcal{G}}(v) \leqslant m$ then $\mathsf{Val}_{\mathcal{G}^n}(v, n) \leqslant m$.*

Proof. By hypothesis and using the memoryless determinacy of [10], there exists a memoryless strategy σ_{Min} for Min in $\mathcal{G}$ such that $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Min}}) \leqslant m$. Let σ^m_{Min} be the strategy in $\mathcal{G}^n$ defined, for all finite play π, vertex v' and $1 \leqslant j \leqslant n$, by

$$\sigma^m_{\mathsf{Min}}(\pi(v', j)) = (\mathtt{in}, \sigma_{\mathsf{Min}}(\mathsf{proj}(\pi)v'), j)\,,$$
$$\sigma^m_{\mathsf{Min}}(\pi(\mathtt{in}, v', j)) = \begin{cases} (v', j) & \text{if } \mathbf{TP}(\pi(\mathtt{in}, v', j)) \geqslant m+1\,, \\ (\mathtt{ex}, v', j) & \text{if } \mathbf{TP}(\pi(\mathtt{in}, v', j)) \leqslant m\,. \end{cases}$$

Intuitively σ^m_{Min} simulates σ_{Min}, and asks to leave the copy when the current total-payoff is less than or equal to m. Notice that, by construction of σ^m_{Min}, $\mathsf{proj}(\pi)$ conforms to σ_{Min}, if π conforms to σ^m_{Min}.

As a first step, if a play π starting in (v, n) and conforming to σ^m_{Min} encounters the target then its value is at most m. Indeed, it is of the form $\pi = \pi'(\mathtt{in}, v', j)(\mathtt{ex}, v', j)\mathtt{t}^\omega$, and since it conforms to σ^m_{Min} we have

$$\mathbf{MCR}(\pi) = \mathbf{TP}(\pi'(\mathtt{in}, v', j)(\mathtt{ex}, v', j)\mathtt{t}) = \mathbf{TP}(\pi'(\mathtt{in}, v', j)) \leqslant m.$$

Then, assume, by contradiction, that there exists a play π starting in (v, n) and conforming to σ^m_{Min}, that does not encounter the target. Then, this means that Min does not ask $n+1$ times the ability to exit in π (since on the $(n+1)$th time that we jump in an exterior vertex, Max is forced to go to the target). In particular, there exists $0 \leqslant j \leqslant n$ such that π is of the form $\pi'(v_1, j)(\mathtt{in}, v_2, j)(v_2, j)(\mathtt{in}, v_3, j)\cdots(v_k, j)(\mathtt{in}, v_k, j)\cdots$. Since for all i, $\sigma^m_{\mathsf{Min}}(\pi'(v_1, j)(\mathtt{in}, v_2, j)\cdots(\mathtt{in}, v_i, j)) = (v_i, j)$, we have that $\mathbf{TP}(\pi'(v_1, j)(\mathtt{in}, v_2, j)\cdots(\mathtt{in}, v_i, j)) \geqslant m+1$. Therefore, since any prefix of $\mathsf{proj}(\pi)$ is the projection of a prefix of π, Lemma 18 shows that $\mathbf{TP}(\mathsf{proj}(\pi)) \geqslant m+1 > m$, which raises a contradiction since $\mathsf{proj}(\pi)$ conforms to σ_{Min} and $\mathsf{Val}(v, \sigma_{\mathsf{Min}}) \leqslant m$. Hence every play that conforms to σ^m_{Min} encounters the target, and, hence, has value at most m. This implies that $\mathsf{Val}_{\mathcal{G}^n}(v, n) \leqslant m$. ◀

We now turn to the other comparison between $\mathsf{Val}_{\mathcal{G}^n}(v, n)$ and $\mathsf{Val}_{\mathcal{G}}(v)$. Since $\mathsf{Val}_{\mathcal{G}}(v)$ can be infinite in case the target is not reachable, we have to be more careful. In particular, we show that $\mathsf{Val}_{\mathcal{G}^n}(v, n) \geqslant \min\big(\mathsf{Val}_{\mathcal{G}}(v), (|V|-1)W+1\big)$ holds for large values of n. In the following, we let $K = |V|(2(|V|-1)W+1)$.

▶ **Lemma 22.** *For all $m \leqslant (|V|-1)W+1$, $k \geqslant K$, and vertex v, if $\mathsf{Val}_{\mathcal{G}}(v) \geqslant m$ then $\mathsf{Val}_{\mathcal{G}^k}(v, k) \geqslant m$.*

Proof. By hypothesis and using the memoryless determinacy of [10], there exists a memoryless strategy σ_{Max} for Max in $\mathcal{G}$ such that $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}}) \geqslant m$. Let σ^m_{Max} be the strategy in $\mathcal{G}^K$ defined, for all finite play π, vertex v' and $1 \leqslant j \leqslant n$, by:

$$\sigma^m_{\mathsf{Max}}(\pi(v, j)) = (\mathtt{in}, \sigma_{\mathsf{Max}}(\mathsf{proj}(\pi)v), j)\,,$$
$$\sigma^m_{\mathsf{Max}}(\pi(\mathtt{ex}, v, j)) = \begin{cases} (v, j-1) & \text{if } \mathbf{TP}(\pi) \leqslant m-1 \text{ and } j > 1\,, \\ \mathtt{t} & \text{otherwise}\,. \end{cases}$$

Intuitively σ^m_{Max} simulates σ_{Max}, and accepts to go to the target when the current total-payoff is greater than or equal to m.

By construction of σ^m_{Max}, if π conforms to σ^m_{Max}, then $\mathsf{proj}(\pi)$ conforms to σ_{Max}. From the structure of the weighted graph, we know that for every play π of $\mathcal{G}^k$, there exists $1 \leqslant j \leqslant k$ such that π is of the form $\pi_k(\mathtt{ex}, v_k, k)\pi_{k-1}(\mathtt{ex}, v_{k-1}, k-1)\cdots\pi_j(\mathtt{ex}, v_j, j)\pi'$ verifying that: there are no occurrences of exterior vertices in $\pi_k, \ldots, \pi_j, \pi'$; for all $\ell \leqslant j$, all vertices in π_ℓ belong to the ℓ-th copy of $\mathcal{G}$; either $\pi' = \mathtt{t}^\omega$ or all vertices of π' belong to the $(j+1)$th copy of $\mathcal{G}$ (in which case, $j < k$).

We now show that, in $\mathcal{G}^k$, $\mathbf{MCR}(\pi) \geqslant m$ for all play π starting in (v, k) and conforming to σ^m_{Max}. There are three cases to consider.

1. If π does not reach the target, then $\mathbf{MCR}(\pi) = +\infty \geqslant m$.
2. If $\pi = \pi_k(\mathtt{ex}, v_k, k)\cdots\pi_j(\mathtt{ex}, v_j, j)\mathtt{t}^\omega$ and $j > 1$ then,

$$\sigma^m_{\mathsf{Max}}(\pi_k(\mathtt{ex}, v_0, k)\cdots\pi_j(\mathtt{ex}, v_j, j)) = \mathtt{t}\,.$$

Thus, using Lemma 18,

$$\begin{aligned} \mathbf{MCR}(\pi) &= \mathbf{TP}(\pi_k(\mathtt{ex}, v_k, k)\cdots\pi_j(\mathtt{ex}, v_j, j)\mathtt{t}) \\ &= \mathbf{TP}(\mathsf{proj}(\pi_k(\mathtt{ex}, v_k, k)\cdots\pi_j(\mathtt{ex}, v_j, j)\mathtt{t})) \\ &\geqslant \mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}}) \geqslant m\,. \end{aligned}$$

3. If $\pi = \pi_k(\mathtt{ex}, v_k, k)\cdots\pi_1(\mathtt{ex}, v_1, 1)\mathtt{t}^\omega$, assume by contradiction that

$$\mathbf{TP}(\pi_k(\mathtt{ex}, v_k, k)\cdots\pi_1) \leqslant m-1\,.$$

Otherwise, we directly obtain $\mathbf{MCR}(\pi) \geqslant m$. Let $v^\star$ be a vertex that occurs at least $N = \lceil K/|V|\rceil = 2(|V|-1)W+1$ times in the sequence $v_1,\dots,v_k$: such a vertex exists, since otherwise $K \leqslant k \leqslant (N-1)|V|$ which contradicts the fact that $(N-1)|V| < K$. Let $j_1 > \cdots > j_N$ be a sequence of indices such that $v_{j_i} = v^\star$ for all i. We give a new decomposition of π:

$$\pi = \pi'_1(\mathtt{ex}, v_{j_1}, j_1)\cdots\pi'_N(\mathtt{ex}, v_{j_N}, j_N)\pi'_{N+1}\,.$$

Since π conforms to σ^m_{Max} and according to the assumption, we have that for all i,

$$\mathbf{TP}(\pi'_1(\mathtt{ex}, v_{j_1}, j_1)\cdots\pi'_i) \leqslant m-1\,.$$

We consider two cases.

a. If there exists π'_i such that $\mathbf{TP}(\pi'_i) \leqslant 0$ then, let $\mathsf{proj}(\pi'_i) = u_1\cdots u_\ell$ with $u_1 = u_\ell = v^\star$, Since π'_i conforms to σ^m_{Max}, $\mathsf{proj}(\pi'_i)$ conforms to σ_{Max}. Therefore the play

$$\widetilde{\pi} = \mathsf{proj}(\pi'_1(\mathtt{ex}, v_{j_1}, j_1)\cdots\pi'_i(\mathtt{ex}, v_{j_i}, j_i))(u_1\cdots u_{\ell-1})^\omega$$

conforms to σ_{Max}. Furthermore, using again Lemma 18,

$$\mathbf{TP}(\widetilde{\pi}) = \liminf_{n\to+\infty}\Big(\mathbf{TP}\big(\pi'_1(\mathtt{ex}, v_{j_1}, i_1)\cdots\pi'_i(\mathtt{ex}, v_{j_i}, j_i)\big) + n\mathbf{TP}(u_1\cdots u_\ell)\Big)$$

and since $\mathbf{TP}(u_1\cdots u_\ell) = \mathbf{TP}(\pi'_i) \leqslant 0$, we have

$$\mathbf{TP}(\widetilde{\pi}) \leqslant \mathbf{TP}(\pi'_1(\mathtt{ex}, v_{j_1}, i_1)\cdots\pi'_i(\mathtt{ex}, v_{j_i}, j_i)) \leqslant m-1.$$

Thus $\widetilde{\pi}$ is a play starting from v that conforms to σ_{Max} but whose total-payoff is strictly less than m, which raises a contradiction.

b. If for all π'_i, $\mathbf{TP}(\pi'_i) \geqslant 1$ (notice that it is implied by $\mathbf{TP}(\pi'_i) > 0$). From Lemma 19, since $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}}) \geqslant m \neq -\infty$, we know that $\mathbf{TP}(\mathsf{proj}(\pi'_0)) \geqslant -(|V|-1)W$. From Lemma 18, $\mathbf{TP}(\pi'_0) \geqslant -(|V|-1)W$. Therefore

$$\begin{aligned}\mathbf{TP}(\pi'_1(\mathtt{ex}, v_{j_1}, i_1)\cdots\pi'_N) &\geqslant -(|V|-1)W + N\\ &= (|V|-1)W + 1 \geqslant m\end{aligned}$$

which contradicts the assumption that

$$\mathbf{TP}(\pi'_1(\mathtt{ex}, v_{j_1}, i_1)\cdots\pi'_N) < m.$$

We have shown that $\mathbf{MCR}(\pi) \geqslant m$ for all play π starting in (v,k) and conforming to σ^m_{Max}, which implies $\mathsf{Val}_{\mathcal{G}^K}((v,k), \sigma^m_{\mathsf{Max}}) \geqslant m$. ◀

From the two previous lemmas, we are ready to relate precisely values in $\mathcal{G}$ and $\mathcal{G}^k$.

Proof of Proposition 2. Let $v \in V$.

- If $\mathsf{Val}_{\mathcal{G}}(v) = -\infty$, then for all m, $\mathsf{Val}_{\mathcal{G}}(v) \leqslant m$. Thus, by Lemma 21, $\mathsf{Val}_{\mathcal{G}^K}((v,K)) \leqslant m$. Therefore $\mathsf{Val}_{\mathcal{G}^K}(v,K) = -\infty$.
- If $\mathsf{Val}_{\mathcal{G}}(v) = m \in [-(|V|-1)W, (|V|-1)W]$. Then, $m \leqslant \mathsf{Val}_{\mathcal{G}}(v) \geqslant m$. Thus, by Lemma 21 and 22, $m \leqslant \mathsf{Val}_{\mathcal{G}^K}((v,K)) \geqslant m$. Therefore $\mathsf{Val}_{\mathcal{G}^K}((v,K)) = m$.
- If $\mathsf{Val}_{\mathcal{G}}(v) = +\infty$, then $\mathsf{Val}_{\mathcal{G}}(v) \geqslant (|V|-1)W+1$. Thus, by Lemma 22, $\mathsf{Val}_{\mathcal{G}^K}((v,K)) \geqslant (|V|-1)W+1$.

◀

H Value iteration algorithm for total-payoff games

This section is devoted to the study of Algorithm 2, in particular in the proof of Theorem 3.

We first define formally the game $\mathcal{G}_Y$ described informally on page 11. To the original total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP}\rangle$ and to every vector $Y \in \mathbb{Z}_\infty^V$, we associate the min-cost reachability game $\mathcal{G}_Y = \langle V', E_Y, \omega_Y, \{\mathtt{t}\}\text{-}\mathbf{MCR}\rangle$ as follows. The sets of vertices are given by

$$V'_{\mathsf{Max}} = V_{\mathsf{Max}} \uplus \{(\mathtt{in}, v) \mid v \in V\} \uplus \{\mathtt{t}\} \quad \text{and} \quad V'_{\mathsf{Min}} = V_{\mathsf{Min}}\,.$$

As in game $\mathcal{G}^j$, vertices of the form $(\mathtt{in}, v)$ are called *interior vertices*. Edges are defined by

$$\begin{aligned} E_Y = \quad & \big\{\big(v, (\mathtt{in}, v')\big) \mid (v, v') \in E\big\} \uplus \big\{\big((\mathtt{in}, v), v\big) \mid v \in V\big\} \\ & \uplus \big\{\big((\mathtt{in}, v), \mathtt{t}\big) \mid v \in V \wedge Y(v) \neq +\infty\big\} \uplus \{(\mathtt{t}, \mathtt{t})\} \end{aligned}$$

while weights of edges are defined, for all $(v, v') \in E$, by

$$\begin{aligned} &\omega_Y\big(v, (\mathtt{in}, v')\big) = \omega_Y(v, v')\,, && \omega_Y\big((\mathtt{in}, v), \mathtt{t}\big) = \max\big(0, Y(v)\big)\,, \\ &\omega_Y\big((\mathtt{in}, v), v\big) = 0. \end{aligned}$$

It is easy to see that lines 3 to 10 are a rewriting of Algorithm 1 in the special case of game $\mathcal{G}_Y$: in particular, neither the target vertex nor interior vertices are explicit, but their behavior is taken into account by the transformation performed in line 3 and the operators min used in the inner computation of lines 6 and 7. Hence, if we define $\mathcal{H}(Y)(v) = \mathsf{Val}_{\mathcal{G}_Y}(v)$ for all $v \in V$, we can say that if inside the main loop, at line 3 the variable Y has value Y, then after line 10, it has value $\mathcal{H}(Y)$.

Notice that the game $\mathcal{G}_Y$ resembles a copy of $\mathcal{G}$ in the game $\mathcal{G}^j$ of the previous section. More, precisely, from the values $(\mathsf{Val}_{\mathcal{G}^j}(v, j))_{v\in V}$ in the jth copy, we can deduce the values in the $(j+1)$th copy by an application of operator $\mathcal{H}$:

$$\big(\mathsf{Val}_{\mathcal{G}^{j+1}}(v, j+1)\big)_{v\in V} = \mathcal{H}\big((\mathsf{Val}_{\mathcal{G}^j}(v, j))_{v\in V}\big)\,.$$

Although the 0th copy is not defined, we abuse the notation and set $\mathsf{Val}_{\mathcal{G}^0}(v, 0) = -\infty$, which still conforms to the above equality. Furthermore, due to the structure of the game $\mathcal{G}^j$ notice that for all $j \leqslant j'$, $\mathsf{Val}_{\mathcal{G}^j}(v, j) = \mathsf{Val}_{\mathcal{G}^{j'}}(v, j)$.

Notice the absence of exterior vertices $(\mathtt{ex}, v', j)$ in game $\mathcal{G}_Y$, replaced by the computation of the maximum between 0 and $X(v')$ on the edge towards the target. Before proving the correctness of Algorithm 2, we prove several interesting properties of operator $\mathcal{H}$.

▶ Proposition 23. $\mathcal{H}$ is a monotonic operator.

Proof. For every vector $Y \in \mathbb{Z}_\infty^V$, let $\mathcal{F}_Y$ be the operator associated with the min-cost reachability game as defined in Section 3, i.e., for all $X \in \mathbb{Z}_\infty^{V'}$, and for all $v_1 \in V'$

$$\mathcal{F}_Y(X)(v_1) = \begin{cases} \max\limits_{v_2 \in E_Y(v_1)} \big(\omega_Y(v_1, v_2) + X(v_2)\big) & \text{if } v \in V'_{\mathsf{Max}} \setminus \{\mathtt{t}\} \\ \min\limits_{v_2 \in E_Y(v_1)} \big(\omega_Y(v_1, v_2) + X(v_2)\big) & \text{if } v \in V_{\mathsf{Min}} \\ 0 & \text{if } v_1 = \mathtt{t}\,. \end{cases}$$

We know from Corollary 11 that $\mathsf{Val}_{\mathcal{G}_Y}$ is the greatest fixed point of $\mathcal{F}_Y$. Consider now two vectors $Y, Y' \in \mathbb{Z}_\infty^V$ such that $Y \preccurlyeq Y'$.

First, notice that for all $X \in \mathbb{Z}_\infty^{V'}$:

$$\mathcal{F}_Y(X) \preccurlyeq \mathcal{F}_{Y'}(X)\,. \tag{5}$$

Indeed, to get the result it suffices to notice that for all $v_1, v_2 \in V'$, $\omega_Y(v_1, v_2) \leqslant \omega_{Y'}(v_1, v_2)$.

Consider then the vector X_0 defined by $X_0(v_1) = +\infty$ for all $v_1 \in V'$. From (5), we have that $\mathcal{F}_Y(X_0) \preccurlyeq \mathcal{F}_{Y'}(X_0)$, then a simple induction shows that for all i, $\mathcal{F}^i_Y(X_0) \preccurlyeq \mathcal{F}^i_{Y'}(X_0)$. Thus, since $\mathsf{Val}_{\mathcal{G}_Y}$ (respectively, $\mathsf{Val}_{\mathcal{G}_{Y'}}$) is the greatest fixed point of $\mathcal{F}_Y$ (respectively, $\mathcal{F}_{Y'}$), we have $\mathsf{Val}_{\mathcal{G}_Y} \preccurlyeq \mathsf{Val}_{\mathcal{G}_{Y'}}$. As a consequence $\mathcal{H}(Y) \preccurlyeq \mathcal{H}(Y')$. ◀

Notice that $\mathcal{H}$ may not be Scott-continuous, as shown in the following example.

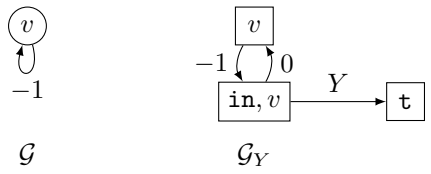

$\mathcal{G}$ $\mathcal{G}_Y$

Figure 6 The games $\mathcal{G}$ and $\mathcal{G}_Y$.

▶ **Example 24.** Recall that, in our setting, a Scott-continuous operator is a mapping $F : \mathbb{Z}_\infty^V \to \mathbb{Z}_\infty^V$ such that for any sequence of vectors $(x_i)_{i\geqslant 0}$ having a limit x_ω, the sequence $(F(x_i))_{i\geqslant 0}$ has a limit equal to $F(x_\omega)$.

We present a total-payoff game whose associated operator $\mathcal{H}$ is not continuous. Let $\mathcal{G}$ be the total-payoff game containing one vertex v of Min and a self loop of weight -1 (as depicted in Fig. 6). For all $Y \in \mathbb{Z}$, in the min-cost reachability game $\mathcal{G}_Y$, v has value $-\infty$, indeed one can take the loop an arbitrary number of times before reaching the target, ensuring a value arbitrary low. Therefore, if we take an increasing sequence $(Y_i)_{i\geqslant 0}$ of integers, $\mathcal{H}(Y_i)(v) = -\infty$ for all i, thus the limit of the sequence $(\mathcal{H}(Y_i))_{i\geqslant 0}$ is $-\infty$. However, the limit of the sequence $(Y_i)_{i\geqslant 0}$ is $+\infty$ and $\mathcal{H}(+\infty)(v) = +\infty$, since the target is not reachable anymore (in case the weight of an edge would be $+\infty$, it is removed in the definition of E_Y). Thus, $\mathcal{H}$ is not Scott-continuous. ◀

In particular, we may not use the Kleene sequence, as we have done for min-cost reachability games, to conclude to the correctness of our algorithm. Anyhow, we will show that the sequence $(Y^j)_{j\geqslant 0}$ indeed converges towards the vector of values of the total-payoff game. We first show that this vector is a pre-fixed point of $\mathcal{H}$ starting with a technical lemma that is useful in the subsequent proof.

▶ **Lemma 25.** *Let σ_{Min} be a strategy for* Min *in $\mathcal{G}$, and $\pi = v_1 \cdots v_i$ a finite play that conforms to σ_{Min}. Then:*

$$\mathbf{TP}(v_1 \cdots v_i) + \mathsf{Val}_{\mathcal{G}}(v_i) \leqslant \mathsf{Val}_{\mathcal{G}}(v_1, \sigma_{\mathsf{Min}})\,.$$

Proof. Let σ_{Max} be an optimal strategy for Min and $v_i v_{i+1} v_{i+2} \cdots$ be the play $\mathsf{Play}(v_i, \sigma_{\mathsf{Max}}, \sigma_{\mathsf{Min}})$. Since σ_{Max} is optimal, $\mathbf{TP}(v_i v_{i+1} v_{i+2} \cdots) \geqslant \mathsf{Val}_{\mathcal{G}}(v_i)$. Furthermore notice that $v_1 \cdots v_i v_{i+1} \cdots$ conforms to the strategy σ_{Min}, therefore $\mathbf{TP}(v_1 v_1 \cdots v_i v_{i+1} \cdots) \leqslant \mathsf{Val}_{\mathcal{G}}(v_1, \sigma_{\mathsf{Min}})$. Thus:

$$\begin{aligned}\mathbf{TP}(v_1 \cdots v_i) + \mathsf{Val}_{\mathcal{G}}(v_i) &\leqslant \mathbf{TP}(v_1 \cdots v_i) + \mathbf{TP}(v_i v_{i+1} \cdots) \\ &\leqslant \mathbf{TP}(v_1 v_1 \cdots v_i v_{i+1} \cdots) \leqslant \mathsf{Val}_{\mathcal{G}}(v_1, \sigma_{\mathsf{Min}})\,.\end{aligned}$$

◀

▶ **Lemma 26.** $\mathsf{Val}_{\mathcal{G}}$ *is a pre-fixed point of $\mathcal{H}$, i.e., $\mathcal{H}(\mathsf{Val}_{\mathcal{G}}) \preccurlyeq \mathsf{Val}_{\mathcal{G}}$.*

Proof. To ease the notations, we denote $\mathsf{Val}_{\mathcal{G}}$ by $Y^\star$ in this proof. To prove this lemma, we just have to show that for all $v_1 \in V$, the value of v_1 in the min-cost reachability game $\mathcal{G}_{Y^\star}$ is at most its value in the original total-payoff game $\mathcal{G}$, i.e.,

$$\mathcal{H}(Y^\star)(v_1) = \mathsf{Val}_{\mathcal{G}_{Y^\star}}(v_1) \leqslant Y^\star(v_1)\,.$$

Let σ_{Min} be a memoryless strategy in $\mathcal{G}$ such that $\mathsf{Val}_{\mathcal{G}}(v_1, \sigma_{\mathsf{Min}}) \leqslant m$ for some $m \in \mathbb{Z} \uplus \{+\infty\}$. And let σ^m_{Min} be a strategy in $\mathcal{G}_X$ defined for all finite play $\pi v'$ with $v' \in V_{\mathsf{Min}}$ by $\sigma^m_{\mathsf{Min}}(\pi v') = (\mathtt{in}, \sigma_{\mathsf{Min}}(v'))$ and for all finite play $\pi(\mathtt{in}, v')$,

$$\sigma^m_{\mathsf{Min}}(\mathtt{in}, v') = \begin{cases} \mathtt{t} & \text{if } \mathbf{TP}(\pi(\mathtt{in}, v')\mathtt{t}) \leqslant m \\ v & \text{otherwise}\,. \end{cases}$$

Notice that, by construction, all plays starting in v_1, conforming to σ^m_{Min} and that reach the target have value at most m. Assume by contradiction that there exists a play $v_1(\mathtt{in}, v_2)v_2(\mathtt{in}, v_3)\cdots \in \mathsf{Play}(v_1, \sigma^m_{\mathsf{Min}})$ that never reaches $\mathtt{t}$. In particular, for all i, $\mathbf{TP}(v_1(\mathtt{in}, v_2)\cdots(\mathtt{in}, v_i)\mathtt{t}) > m$.

Again by construction, $v_1 v_2 \cdots$ is a play in $\mathcal{G}$ that conforms to σ_{Min} and for all i, $\mathbf{TP}(v_1(\mathtt{in}, v_2)\cdots(\mathtt{in}, v_i)) = \mathbf{TP}(v_1 \cdots v_i)$.

- If there exists $i \geqslant 2$ such that $Y^\star(v_i) = \mathsf{Val}_{\mathcal{G}}(v_i) \geqslant 0$, then

$$\begin{aligned}\mathbf{TP}(v_1(\mathtt{in}, v_2)\cdots(\mathtt{in}, v_i)\mathtt{t}) &= \mathbf{TP}(v_1(\mathtt{in}, v_2)\cdots(\mathtt{in}, v_i)) + \omega_X((\mathtt{in}, v_i), \mathtt{t}) \\ &= \mathbf{TP}(v_1 \cdots v_i) + \max(0, Y^\star(v_i)) \\ &= \mathbf{TP}(v_1 \cdots v_i) + Y^\star(v_i) \\ &\leqslant \mathsf{Val}(v_1, \sigma_{\mathsf{Min}}) \qquad \text{(from Lemma 25)} \\ &\leqslant m\end{aligned}$$

which raises a contradiction.

- If for all $i \geqslant 2$, $Y^\star(v_i) = \mathsf{Val}_{\mathcal{G}}(v_i) < 0$ then for all $i \geqslant 2$,

$$\mathbf{TP}(v_1(\mathtt{in}, v_2)\cdots(\mathtt{in}, v_i)\mathtt{t}) = \mathbf{TP}(v_1(\mathtt{in}, v_2)\cdots(\mathtt{in}, v_i)) = \mathbf{TP}(v_1\cdots v_i) > m\,.$$

Thus $\mathbf{TP}(v_1v_2\cdots) > m$, which contradicts the fact that $v_1v_2\cdots$ conforms to σ_{Min} and $\mathsf{Val}_{\mathcal{G}}(v_1, \sigma_{\mathsf{Min}}) \leqslant m$.

Thus $\mathsf{Val}_{\mathcal{G}_{Y^\star}}(v_1, \sigma_{\mathsf{Min}}) \leqslant m$. As a consequence, $\mathsf{Val}_{\mathcal{G}_{Y^\star}}(v_1) \leqslant \mathsf{Val}_{\mathcal{G}}(v_1)$. ◀

▶ Remark. Even if it is not necessary for the proof of Theorem 3, we can show that $\mathsf{Val}_{\mathcal{G}}$ is the least pre-fixed point of $\mathcal{H}$. Notice that, by monotonicity of $\mathcal{H}$, this directly implies that $\mathsf{Val}_{\mathcal{G}}$ is the least fixed point of $\mathcal{H}$. The proof that $\mathsf{Val}_{\mathcal{G}}$ is the least pre-fixed point of $\mathcal{H}$ amounts to better understand the convergence of the sequence $(\mathsf{Val}_{\mathcal{G}^j}(v,j))_{j\geqslant 0}$ (remember that we set $\mathsf{Val}_{\mathcal{G}^0}(v,0) = -\infty$ for all v). Indeed, Proposition 2 already shows that for vertices v such that $\mathsf{Val}_{\mathcal{G}}(v) < +\infty$, $(\mathsf{Val}_{\mathcal{G}^j}(v,j))_{j\geqslant 1}$ converges towards $\mathsf{Val}_{\mathcal{G}}(v)$. It is also the case for vertices v such that $\mathsf{Val}_{\mathcal{G}}(v) = +\infty$ as we show in Lemma 27 below. Then, consider a pre-fixed point Y of $\mathcal{H}$, i.e., $\mathcal{H}(Y) \preccurlyeq Y$. Since $\mathsf{Val}_{\mathcal{G}^0}(v,0) = -\infty \leqslant Y(v)$ and $\mathcal{H}$ is monotonous, we prove by immediate induction that $\mathsf{Val}_{\mathcal{G}^j}(v,j) \leqslant Y(v)$ for all v and $j \geqslant 0$: indeed, if $\mathsf{Val}_{\mathcal{G}^j}(v,j) \leqslant Y(v)$ for all v, we have

$$(\mathsf{Val}_{\mathcal{G}^{j+1}}(v,j+1))_{v\in V} = \mathcal{H}((\mathsf{Val}_{\mathcal{G}^j}(v,j))_{v\in V}) \preccurlyeq \mathcal{H}(Y) \preccurlyeq Y\,.$$

This implies that $\mathsf{Val}_{\mathcal{G}} \preccurlyeq Y$, showing that $\mathsf{Val}_{\mathcal{G}}$ is indeed the least pre-fixed point of $\mathcal{H}$, and hence the least fixed point of $\mathcal{H}$, by the above reasoning.

Before continuing the proof of Theorem 3, we show the result used in the previous remark.

▶ **Lemma 27.** *Let $v \in V$ such that $\mathsf{Val}_{\mathcal{G}}(v) = +\infty$, and σ_{Max} a memoryless strategy for Max in $\mathcal{G}$ such that $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}}) = +\infty$. Then the following holds:*

(i) For every finite play $v_1\cdots v_k$ conforming to σ_{Max} starting in $v_1 = v$, if there exists $i < j$ such that $v_i = v_j$ then $\mathbf{TP}(v_i\cdots v_j) \geqslant 1$.
(ii) For every $m \in \mathbb{N}$, $k \geqslant m|V| + 1$ and $v_1\cdots v_k$ a finite play conforming to σ_{Max} and starting in $v_1 = v$, $\mathbf{TP}(v_1\cdots v_k) \geqslant m - (|V|-1)W$.
(iii) For all $m \in \mathbb{N}$ and $k \geqslant (m + (|V|-1)W)|V| + 1$, $\mathsf{Val}_{\mathcal{G}^k}(v,k) \geqslant m$.
(iv) $\lim_{j\to\infty} \mathsf{Val}_{\mathcal{G}^j}(v,j) = +\infty$.

Proof. We prove (i) by contradiction. Therefore, assume that $\mathbf{TP}(v_i\cdots v_j) \leqslant 0$, then $\pi = v_1\cdots v_{i-1}(v_i\cdots v_{j-1})^\omega$ comforms to σ_{Max} and $\mathbf{TP}(\pi) \leqslant \mathbf{TP}(v_1\cdots v_{i-1}) < +\infty$, which contradicts the fact that $\mathsf{Val}_{\mathcal{G}}(v, \sigma_{\mathsf{Max}}) = +\infty$.

We prove (ii) by induction on m. The base case is straightforward. For the inductive case, let $m > 0$, $k \geqslant m|V| + 1$ and $v_1\cdots v_k$ a finite play conforming to σ_{Max} and starting in $v_1 = v$. Since $k \geqslant m|V| + 1 \geqslant |V| + 1$ there exists $i < j \leqslant |V| + 1$ such that $v_i = v_j$. Thus:

$$\begin{aligned}\mathbf{TP}(v_1\cdots v_k) &= \mathbf{TP}(v_1\cdots v_i) + \mathbf{TP}(v_i\cdots v_j) + \mathbf{TP}(v_j\cdots v_k)\\ &= \mathbf{TP}(v_1\cdots v_iv_{j+1}\cdots v_k) + \mathbf{TP}(v_i\cdots v_j)\\ &\geqslant \mathbf{TP}(v_1\cdots v_iv_{j+1}\cdots v_k) + 1\,. && \text{from } (i)\end{aligned}$$

By induction hypothesis, as $v_1\cdots v_iv_{j+1}\cdots v_k$ conforms to σ_{Max} and has length at least $(m-1)|V| + 1$ (because $i < j \leqslant |V| + 1$, implying that $j - i \leqslant |V|$), we have $\mathbf{TP}(v_1\cdots v_iv_{j+1}\cdots v_k) \geqslant m - 1 - (|V|-1)W$, thus $\mathbf{TP}(v_1\cdots v_k) \geqslant m - (|V|-1)W$.

To prove (iii), let σ'_{Max} be a strategy of Max in $\mathcal{G}^k$ defined by $\sigma'_{\mathsf{Max}}(v,j) = (\mathtt{in}, \sigma_{\mathsf{Max}}(v), j)$ and $\sigma'_{\mathsf{Max}}(\mathtt{ex}, v, j) = (v, j-1)$ for all $v \in V$ and $j \geqslant k$. Let π be a play starting in (v,k) and conforming to σ'_{Max}. If π does not reach $\mathtt{t}$, then $\mathbf{MCR}(\pi) = +\infty \geqslant m$. If π reaches the target then $\mathsf{proj}(\pi)$ is of the form $v_1\cdots v_\ell\mathtt{t}^\omega$, with $\mathbf{MCR}(\pi) = \mathbf{TP}(v_1\cdots v_\ell)$. It is clear by construction of σ'_{Max} that $v_1\cdots v_\ell$ is a finite play of $\mathcal{G}$ that conforms to σ_{Max}. Furthermore, $\ell \geqslant k \geqslant (m + (|V|-1)W)|V| + 1$ thus, from (ii), we have that $\mathbf{TP}(v_1\cdots v_\ell) \geqslant m$. This implies $\mathbf{MCR}(\pi) \geqslant m$. Hence, every play in $\mathcal{G}^k$ conforming to σ'_{Max} and starting in (v,k) has a value at least m, which means that $\mathsf{Val}_{\mathcal{G}^k}(v,k) \geqslant m$.

Item (iv) is then a direct consequence of (iii). ◀

We are now ready to state and prove the inductive invariant allowing us to show the correctness of Algorithm 2.

▶ **Lemma 28.** *Before the j-th iteration of the external loop of Algorithm 2, we have $\mathsf{Val}_{\mathcal{G}^j}(v,j) \leqslant Y^j(v) \leqslant \mathsf{Val}_{\mathcal{G}}(v)$ for all vertices $v \in V$.*

Proof. For $j = 0$, we have $Y^0(v) = -\infty = \mathsf{Val}_{\mathcal{G}^0}(v, 0)$ for all vertex $v \in V$. Suppose then that the invariant holds for $j \geqslant 0$. We know that $\mathsf{Val}_{\mathcal{G}^{j+1}}(v, j+1) = \mathcal{H}((\mathsf{Val}_{\mathcal{G}^j}(v', j))_{v' \in V})$. Moreover, after the assignment of line 10, by definition of $\mathcal{H}$, variable Y contains $\mathcal{H}(Y^j)$. The operation performed on line 11 only increases the values of vector Y, so that at the end of the jth iteration, we have $\mathcal{H}(Y^j) \preccurlyeq Y^{j+1}$. Since $\mathcal{H}$ is monotonous, and by the invariant at step j, we obtain

$$\mathsf{Val}_{\mathcal{G}^{j+1}}(v, j+1) = \mathcal{H}((\mathsf{Val}_{\mathcal{G}^j}(v', j))_{v' \in V}) \leqslant \mathcal{H}(Y^j) \leqslant Y^{j+1} .$$

Moreover, using again the monotony of $\mathcal{H}$ and Lemma 26, we have

$$\mathcal{H}(Y^j) \preccurlyeq \mathcal{H}(\mathsf{Val}_{\mathcal{G}}) \preccurlyeq \mathsf{Val}_{\mathcal{G}} .$$

A closer look at line 11 shows that $\mathcal{H}(Y^j)$ and Y^{j+1} coincide over vertices v such that $\mathcal{H}(Y^j)(v) \leqslant (|V|-1)W$, and otherwise $Y^{j+1}(v) = +\infty$. Hence, if $\mathcal{H}(Y^j)(v) \leqslant (|V|-1)W$, we directly obtain $Y^{j+1}(v) = \mathcal{H}(Y^j)(v) \leqslant \mathsf{Val}_{\mathcal{G}}(v)$. Otherwise, we know that $\mathsf{Val}_{\mathcal{G}}(v) > (|V|-1)W$. By Corollary 20, we know that $\mathsf{Val}_{\mathcal{G}}(v) = +\infty$, so that $Y^{j+1}(v) = +\infty = \mathsf{Val}_{\mathcal{G}}(v)$. In the overall, we have proved

$$\mathsf{Val}_{\mathcal{G}^{j+1}}(v, j+1) \leqslant Y^{j+1}(v) \leqslant \mathsf{Val}_{\mathcal{G}}(v) \tag{6}$$

◀

We are now able to prove the correction of the algorithm.

Proof of Theorem 3. For $j = K$ (remember that K was defined in the previous section), the invariant of Lemma 28 becomes

$$\mathsf{Val}_{\mathcal{G}^K}(v, K) \leqslant Y^K(v) \leqslant \mathsf{Val}_{\mathcal{G}}(v)$$

for all vertices $v \in V$. Notice that the iteration may have stopped before iteration K, in which case the sequence $(Y^j)_{j \geqslant 0}$ may be considered as stationary. In case $\mathsf{Val}_{\mathcal{G}}(v) \neq +\infty$, Proposition 2 proves that $\mathsf{Val}_{\mathcal{G}^K}(v, K) = \mathsf{Val}_{\mathcal{G}}(v)$, so that we have $Y^K(v) = \mathsf{Val}_{\mathcal{G}}(v)$. In case $\mathsf{Val}_{\mathcal{G}}(v) = +\infty$, Proposition 2 shows that $\mathsf{Val}_{\mathcal{G}^K}(v, K) > (|V|-1)W$: by the operation performed at line 11, we obtain that $Y^K(v) = +\infty = \mathsf{Val}_{\mathcal{G}}(v)$.

Hence, $K = |V|(2(|V|-1)W + 1)$ is an upper bound on the number of iterations before convergence of Algorithm 2, and moreover, at the convergence, the algorithm outputs the vector of optimal values of the total-payoff game. ◀

I An example of parametric total-payoff game

We depict in Fig. 7 a weighted graph parametrized with the number n of *layers*, as well as the greatest weight $W > 0$. For both the min-cost reachability objective (with $\mathtt{t}$ the target) and the total-payoff objectives, the values of the vertices are as follows: vertices v_{3k+1} and v_{3k+2} ($k \in \{0, \dots, n-1\}$) have value 0, whereas vertices v_{3k} ($k \in \{1, \dots, n\}$) have value W. In our add-on prototype of PRISM, we model the min-cost objective with $\langle\langle \mathsf{Max} \rangle\rangle \mathsf{R}^{\max=?}[\mathsf{F}^{\infty}\mathtt{t}]$ for Max and $\langle\langle \mathsf{Min} \rangle\rangle \mathsf{R}^{\min=?}[\mathsf{F}^{\infty}\mathtt{t}]$ for Min, whereas total-payoff objectives are modelled by $\langle\langle \mathsf{Max} \rangle\rangle \mathsf{R}^{\max=?}[\mathsf{F}^{c}\bot]$ and $\langle\langle \mathsf{Min} \rangle\rangle \mathsf{R}^{\min=?}[\mathsf{F}^{c}\bot]$.

We present in the following table the time for resolution (in seconds), the number of iterations in the external loop, and the total number of iterations in the internal loops for the total-payoff resolution, for various values of parameters W and n:

$W \backslash n$	100	200	300	400	500
50	0.52 / 151 / 12603	1.90 / 251 / 22703	3.84 / 351 / 32803	6.05 / 451 / 42903	9.83 / 551 / 53003
100	1.00 / 201 / 30103	3.48 / 301 / 50203	8.64 / 401 / 70303	13.53 / 501 / 90403	22.64 / 601 / 110503
150	1.89 / 251 / 52603	6.02 / 351 / 82703	12.88 / 451 / 112803	22.13 / 551 / 142903	34.16 / 651 / 173003
200	2.96 / 301 / 80103	9.62 / 401 / 120203	18.33 / 501 / 160303	30.42 / 601 / 200403	45.64 / 701 / 240503
250	3.92 / 351 / 112603	13.28 / 451 / 162703	25.18 / 551 / 212803	46.23 / 651 / 262903	71.51 / 751 / 313003

Notice that due to the very little memory consumption of the algorithm, there is no risk of running out of memory. However, the execution time can become very large. For instance, in case $W = 500$ and $n = 1000$, the execution time becomes $536s$ whereas the total number of iterations in the internal loop is greater than a million.

On this example, with $n+1$ components, the acceleration heuristics presented in details in Appendix J gives excellent results. Indeed, by combining both heuristics, we obtain the following results:

$W \backslash n$	100	200	300	400	500
50	0.01 / 402 / 1404	0.08 / 802 / 2804	0.22 / 1202 / 4204	0.38 / 1602 / 5604	0.42 / 2002 / 7004
100	0.02 / 402 / 1404	0.09 / 802 / 2804	0.19 / 1202 / 4204	0.33 / 1602 / 5604	0.40 / 2002 / 7004
150	0.03 / 402 / 1404	0.09 / 802 / 2804	0.18 / 1202 / 4204	0.29 / 1602 / 5604	0.47 / 2002 / 7004
200	0.02 / 402 / 1404	0.07 / 802 / 2804	0.16 / 1202 / 4204	0.23 / 1602 / 5604	0.47 / 2002 / 7004
250	0.01 / 402 / 1404	0.07 / 802 / 2804	0.17 / 1202 / 4204	0.29 / 1602 / 5604	0.48 / 2002 / 7004

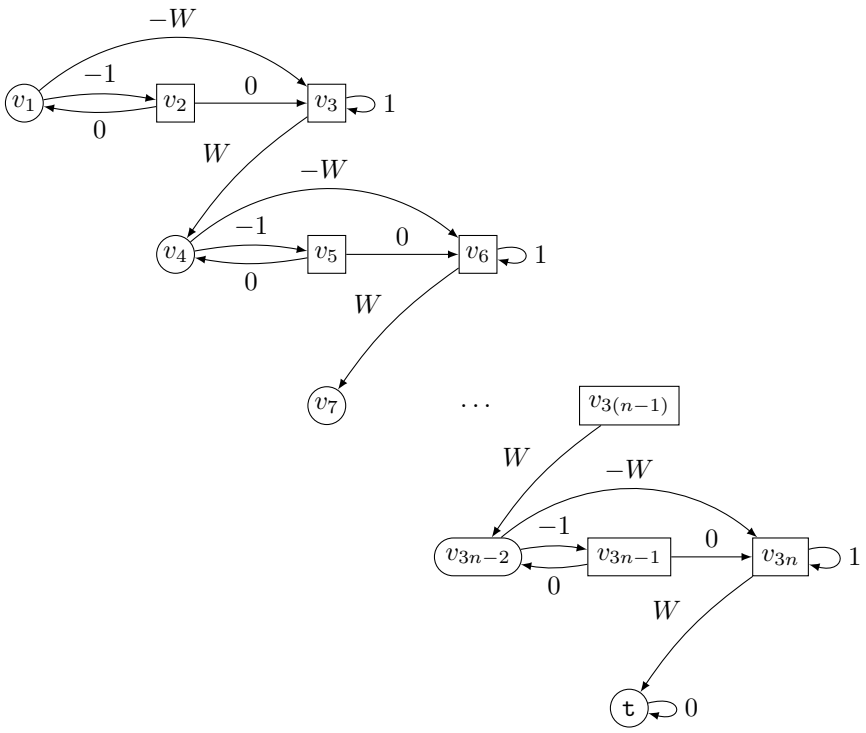

Figure 7 Parametric weighted graph

Notice that the number of iterations in both internal and external loops do no longer depend on the choice of parameter W, as well as the execution time. With respect to the execution time, the decrease from the case without acceleration is even larger, since the updates of vector X inside the inner loop need only to be performed on the vertices of the current component. For large instances, the execution time may again become very large, but in case $W = 500$ (as previously said, this value is independent of the result) and $n = 1000$, it shrinks to $2.3s$ whereas the total number of iterations in the internal loop becomes 14004, i.e., 5 orders of magnitude less than for the algorithm without acceleration heuristics.

J Acceleration heuristics in MCR games

Algorithm 4 and 5 are enhanced versions of Algorithm 1 and 2 respectively, that apply the acceleration heuristics described at the end of Section 5.

Algorithm 4: Accelerated value iteration algorithm for min-cost reachability games

Input: min-cost reachability game $\langle V, E, \omega, \{\mathtt{t}\}\text{-}\mathbf{MCR}\rangle$, SCC-decomposition $\mathtt{c}\colon V \to \{0, 1, \ldots, p\}$ and an oracle $\mathcal{O}(q, v)$ outputting sets $(S_v)_{v \in \mathtt{c}^{-1}(q)}$

$\mathsf{X}(\mathtt{t}) := 0$
for $q = 1$ **to** p **do**
  $(S_v)_{v \in \mathtt{c}^{-1}(q)} := \mathcal{O}(q, \mathsf{X})$ `/* Use of the oracle */`
  foreach $v \in \mathtt{c}^{-1}(q)$ **do** $\mathsf{X}(v) := \max S_v$
  repeat
    $\mathsf{X}_{pre} := \mathsf{X}$
    foreach $v \in \mathtt{c}^{-1}(q) \cap V_{\mathsf{Max}}$ **do**
      $\mathsf{X}(v) := \max_{v' \in E(v)} \big(\omega(v, v') + \mathsf{X}_{pre}(v')\big)$
    foreach $v \in \mathtt{c}^{-1}(q) \cap V_{\mathsf{Min}}$ **do**
      $\mathsf{X}(v) := \min_{v' \in E(v)} \big(\omega(v, v') + \mathsf{X}_{pre}(v')\big)$
    foreach $v \in \mathtt{c}^{-1}(q)$ **do** $\mathsf{X}(v) := \max\big(S_v \cap [-\infty, \mathsf{X}(v)]\big)$
  until $\mathsf{X} = \mathsf{X}_{pre}$
return X

Algorithm 5: Accelerated value iteration algorithm for total-payoff games

Input: Total-payoff game $\mathcal{G} = \langle V, E, \omega, \mathbf{TP}\rangle$, SCC-decomposition $\mathsf{c}\colon V \to \{0, 1, \dots, p\}$ and an oracle $\mathcal{O}(q, v)$ outputting sets $(S_v)_{v\in\mathsf{c}^{-1}(q)}$

foreach $v \in V$ **do** $\mathsf{Y}(v) := -\infty$
for $q = 0$ **to** p **do**
  $(S_v)_{v\in\mathsf{c}^{-1}(q)} := \mathcal{O}(q, \mathsf{Y})$ `/* Use of the oracle */`
  repeat
    foreach $v \in \mathsf{c}^{-1}(q)$ **do**
      $\mathsf{Y}_{pre}(v) := \mathsf{Y}(v)$; $\mathsf{Y}(v) := \max(0, \mathsf{Y}(v))$; $\mathsf{X}(v) := \max S_v$
    repeat
      $\mathsf{X}_{pre} := \mathsf{X}$
      foreach $v \in V_{\mathsf{Max}} \cap \mathsf{c}^{-1}(q)$ **do**
        $\mathsf{X}(v) := \max_{v'\in E(v)} \left[\omega(v, v') + \min(\mathsf{X}_{pre}(v'), \mathsf{Y}(v'))\right]$
      foreach $v \in V_{\mathsf{Min}} \cap \mathsf{c}^{-1}(q)$ **do**
        $\mathsf{X}(v) := \min_{v'\in E(v)} \left[\omega(v, v') + \min(\mathsf{X}_{pre}(v'), \mathsf{Y}(v'))\right]$
      foreach $v \in \mathsf{c}^{-1}(q)$ **do** $\mathsf{X}(v) := \max\left(S_v \cap [-\infty, \mathsf{X}(v)]\right)$
    until $\mathsf{X} = \mathsf{X}_{pre}$
    $\mathsf{Y} := \mathsf{X}$
    foreach $v \in V$ such that $\mathsf{Y}(v) > (|V| - 1)W$ **do** $\mathsf{Y}(v) := +\infty$
  until $\mathsf{Y} = \mathsf{Y}_{pre}$
return Y